



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
31.10.2018 Bulletin 2018/44

(21) Application number: **18164672.0**

(22) Date of filing: **28.03.2018**

(51) Int Cl.:
H04L 12/713 (2013.01) **H04L 12/46** (2006.01)
H04W 48/18 (2009.01) **H04W 84/04** (2009.01)
H04W 84/06 (2009.01) **H04L 12/723** (2013.01)

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR
Designated Extension States:
BA ME
Designated Validation States:
KH MA MD TN

(30) Priority: **24.04.2017 US 201715495631**

(71) Applicant: **The Boeing Company**
Chicago, IL 60606-2016 (US)

(72) Inventor: **SCOTT, James P.**
Chicago, Illinois 60606-2016 (US)

(74) Representative: **Boult Wade Tennant LLP**
Verulam Gardens
70 Gray's Inn Road
London WC1X 8BT (GB)

(54) **SYSTEM AND METHOD FOR FEDERATED NETWORK TRAFFIC PROCESSING**

(57) A system includes a base station including a first MODEM of a first MODEM pool configured to generate a first data packet (including first data and a first VRF instance identifier) based on a first RF signal and a second MODEM of a second MODEM pool configured to generate a second data packet (including second data and a second VRF instance identifier) based on a second RF signal. The system includes a network device coupled

to the base station and configured to receive the first data packet, transmit a first packet (including the first data and a first header including a first indicator associated with the first VRF instance) to a first device via a network, receive the second data packet, and transmit a second packet (including the second data and a second header including a second indicator associated with the second VRF instance) to a second device.




FIG. 1

Description

FIELD OF THE DISCLOSURE

[0001] The present disclosure is generally related to processing network traffic.

BACKGROUND

[0002] Cellular network providers, also referred to as Tier-1 mobile network operators (MNO), are able to use cellular systems to provide telephone, data, and other services to subscribers via wireless devices, such as mobile phones, tablet computers, etc. A cellular network provider has a designated frequency spectrum (e.g., frequency range) from which the network bandwidth (measured in Hertz (Hz)) and capacity (measured in bits per second (bps)) used in providing cellular services to the subscribers is derived. This bandwidth and capacity may be larger than needed to provide the services offered by the cellular network provider. In order to increase revenue, the cellular network provider may sell a portion of the bandwidth, the capacity, and services to Tier-2 operators, also known as mobile virtual network operators (MVNOs) or virtual network operators (VNOs). In this model, the Tier-1 operator is referred to as a Value-Added Reseller (VAR) that resells bandwidth, capacity, and services that would otherwise be "stranded" (i.e., not used to provide revenue to the VAR). The cellular system is then used to provide access to designated cellular or network services to subscribers of the Tier-2 operators. Although the cellular system is able to process traffic for different Tier-2 operators, in certain systems, the traffic associated with a first Tier-2 operator is not isolated from traffic associated with a second Tier-2 operator, which can pose a potential security risk to the Tier-2 operators (e.g., a risk that traffic associated with a particular Tier-2 operator is processed with and provided to a device associated with a different Tier-2 operator).

SUMMARY

[0003] In a particular implementation, a system includes a base station that includes a first MODEM of a first MODEM pool. The first MODEM is configured to generate a first data packet based on a first radio-frequency (RF) signal. The first data packet includes first data and a first virtual routing and forwarding (VRF) instance identifier. The base station also includes a second MODEM of a second MODEM pool. The second MODEM is configured to generate a second data packet based on a second RF signal. The second data packet includes second data and a second VRF instance identifier. The system also includes a network device coupled to the base station and configured to receive the first data packet and to transmit a first packet to a first device via a network. The first packet includes the first data and a first header that includes a first indicator associated with a first VRF

instance that corresponds to the first VRF instance identifier. The network device is further configured to receive the second data packet and to transmit a second packet to a second device via the network. The second packet includes the second data and a second header that includes a second indicator associated with a second VRF instance that corresponds to the second VRF instance identifier. In a particular implementation, each VRF instance is associated with a different "closed user group (CUG)" (e.g., a different Tier-2 operator or service provider, such as a MVNO, a VNO, or other government or enterprise customer looking to outsource physical infrastructure of their private enterprise network(s)), and encapsulating and processing data associated with a particular VRF instance differently than data associated with other VRF instances enables isolation of traffic associated with the particular VRF instance (e.g., a particular user) from traffic associated with other VRF instances (e.g., other CUGs).

[0004] In another particular implementation, a satellite includes a payload including one or more antennas. The satellite includes a memory configured to store channel mapping data that indicates a mapping of at least one radio frequency (RF) channel to a different RF channel. The satellite further includes a processor configured to initiate transmission of a second RF signal from the one or more antennas to a first device (e.g., a subscriber device) based on the channel mapping data and based on a first service provider associated with the first device. The second RF signal is based on a first RF signal received at the one or more antennas from radio frequency equipment (RFE). The first RF signal corresponds to a first RF channel, and the second RF signal corresponds to a second RF channel that is different from the first RF channel. By mapping service providers (e.g., CUGs, such as Tier-2 operators) to different RF channels, the satellite enables traffic associated with different CUGs to be communicated using different RF channels, thereby isolating traffic on a per-CUG (e.g., per-service provider, per-MVNO, or per-government or enterprise entity) basis.

[0005] In another particular implementation, a method includes receiving a first data packet at a network device from a base station, transmitting a first packet from the network device to a first device via a network, receiving a second data packet at the network device from the base station, and transmitting a second packet from the network device to a second device via the network. The first data packet includes first data and a first VRF instance identifier. The first packet includes the first data and a first header that includes a first indicator associated with a first VRF instance. The second data packet includes second data and a second VRF instance identifier. The second packet includes the second data and a second header that includes a second indicator associated with a second VRF instance.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006]

FIG. 1 is a block diagram that illustrates a first example of an implementation of system that processes network traffic associated with different virtual routing and forwarding (VRF) instances;

FIG. 2 is a block diagram that illustrates a second example of an implementation of system that processes network traffic associated with different VRF instances;

FIG. 3 is a ladder diagram that illustrates a first example of network traffic communicated via the systems of FIGs. 1 and 2;

FIG. 4 is a ladder diagram that illustrates a second example of network traffic communicated via the systems of FIGs. 1 and 2;

FIG. 5 is a flow chart of an example of a method of processing network traffic received from a device based on a service provider associated with the device;

FIG. 6 is a flow chart of an example of a method of processing network traffic to be transmitted to a device based on a service provider associated with the device; and

FIG. 7 is a flow chart of an example of a method of processing network traffic associated with different virtual routing and forwarding (VRF) instances.

DETAILED DESCRIPTION

[0007] Particular implementations are described with reference to the drawings. In the description, common features are designated by common reference numbers throughout the drawings. As used herein, various terminology is used for the purpose of describing particular implementations only and is not intended to be limiting. For example, the singular forms "a," "an," and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It may be further understood that the terms "comprise," "comprises," and "comprising" may be used interchangeably with "include," "includes," or "including." Additionally, it will be understood that the term "wherein" may be used interchangeably with "where." As used herein, "exemplary" may indicate an example, an implementation, and/or an aspect, and should not be construed as limiting or as indicating a preference or a preferred implementation. As used herein, an ordinal term (e.g., "first," "second," "third," etc.) used to modify an element, such as a structure, a component, an operation, etc., does not by itself indicate any priority or order of the element with respect to another element, but rather merely distinguishes the element from another element having a same name (but for use of the ordinal term). As used herein, the term "set" refers to a grouping of one or more elements, and the term "plurality" refers to multiple elements.

[0008] In the present disclosure, terms such as "determining," "calculating," "shifting," "adjusting," etc. may be used to describe how one or more operations are performed. It should be noted that such terms are not to be construed as limiting and other techniques may be utilized to perform similar operations. Additionally, as referred to herein, "generating," "calculating," "using," "selecting," "accessing," and "determining" may be used interchangeably. For example, "generating," "calculating," or "determining" a parameter (or a signal) may refer to actively generating, calculating, or determining the parameter (or the signal) or may refer to using, selecting, or accessing the parameter (or signal) that is already generated, such as by another component or device. As used herein, "coupled" may include "communicatively coupled," "electrically coupled," or "physically coupled," and may also (or alternatively) include any combinations thereof. Two devices (or components) may be coupled (e.g., communicatively coupled, electrically coupled, or physically coupled) directly or indirectly via one or more other devices, components, wires, buses, networks (e.g., a wired network, a wireless network, or a combination thereof), etc. Two devices (or components) that are electrically coupled may be included in the same device or in different devices and may be connected via electronics, one or more connectors, or inductive coupling, as illustrative, non-limiting examples. In some implementations, two devices (or components) that are communicatively coupled, such as in electrical communication, may send and receive electrical signals (digital signals or analog signals) directly or indirectly, such as via one or more wires, buses, networks, etc. As used herein, "directly coupled" may include two devices that are coupled (e.g., communicatively coupled, electrically coupled, or physically coupled) without intervening components.

[0009] Implementations disclosed herein are directed to systems and methods for providing network traffic isolation on a per-user basis within a cellular system. In some implementations, the systems and methods described herein are described with reference to federated networks. In other implementations, the systems and methods described herein apply to other types of networks. A cellular system includes multiple components, such as network devices (e.g., "core network" devices), a base station, a beamforming controller, radio-frequency (RF) equipment (RFE), a channelizer, or other components. An operator (e.g., a Tier-1 operator) of the cellular system may sell access to the cellular system to other "CUGs" to provide various network services. As used herein, "CUGs" refers to Tier-2 operators (e.g., service providers) that pay for access to equipment associated with a Tier-1 operator. For convenience, a member of a CUG is referred to as a user. Tier-2 operators include mobile virtual network operators (MVNOs) and virtual network operators (VNOs), as non-limiting examples. A person who subscribes to a service provided by a Tier-2 operator is referred to herein as a subscriber. By isolating network traffic associated with each user

(e.g., each Tier-2 operator), the systems and methods described improve security of data processed by a cellular system. Additionally, flexibility of the cellular system is increased because data associated with different CUGs is able to be processed according to different rules or criteria (instead of a set of overall rules or criteria for all CUGs).

[0010] To implement per-user (e.g., per-Tier-2 operator or per-service provider) network traffic isolation, traffic associated with different CUGs is encapsulated differently and is routed and forwarded through the cellular system using different routing information. In a particular implementation, a base station and a network device (e.g., a mobile management entity, a packet gateway such as a Packet Data Network Gateway (PGW), a serving gateway (SGW), etc.) are configured to encapsulate traffic associated with different CUGs (e.g., different MVNOs) using different identifiers. To illustrate, in a particular implementation, the network device maintains different virtual routing and forwarding (VRF) instances for each user, which includes maintaining different routing information base (RIB) data for each user. Due to the different RIB data, traffic associated with a first user may traverse a different path through an external network than traffic associated with a second user.

[0011] In a particular implementation, traffic associated with different CUGs is processed by different devices (or different components) within the cellular system. To illustrate, the base station includes multiple MODEM pools (e.g., physical MODEM pools, virtual MODEM pools, or logical MODEM pools) that include one or more MODEMs, and MODEMs of each MODEM pool are configured to process data packets associated with a corresponding user. As a particular example, one or more MODEMs of a first MODEM pool process data packets associated with a first user (as indicated by a VRF instance identifier in the data packets), and one or more MODEMs of a second MODEM pool process data packets associated with a second user.

[0012] In addition to isolating the network traffic through networks and the cellular system by user (e.g., by Tier-2 operator or service provider), the cellular system may enable different CUGs to access the cellular system via different portions of a RF spectrum (e.g., different RF channels or portions thereof). In a particular implementation, a channelizer on-board a satellite is configured to receive RF signals from RF circuitry coupled to the base station and to perform "beam-to-beam" switching (e.g., RF signal switching) such that a RF signal received from the RF circuitry and corresponding to traffic associated with the first service provider is used to generate a second RF signal that corresponds to a first portion of the RF spectrum associated with the first service provider. Similarly, a third RF signal corresponding to traffic associated with the second service provider is used to generate a fourth RF signal that corresponds to a second portion of the RF spectrum associated with the second service provider. The second RF signal and the

fourth RF signal are provided to corresponding user devices associated with (e.g., subscribing to services of) the first service provider and the second service provider, respectively. In a particular implementation, the RF signals are deployed in an overlay pattern at the same location such that multiple devices at the same location are able to access the cellular system via different portions of the RF spectrum (e.g., via different RF channels or portions thereof). The RF signals are mapped to corresponding VRF instances (e.g., to corresponding CUGs) and are converted into data packets, for which processing is isolated by VRF instance, as described herein.

[0013] By isolating network traffic associated with different CUGs (e.g., different service providers, government customers, or enterprise customers), the system described herein improves data security as compared to other systems. For example, data associated with different CUGs is encapsulated and processed using different VRF instance identifiers, thereby separating (e.g., isolating) data associated with a particular user from data associated with other CUGs. Additionally, communications between the cellular system and subscriber devices are performed using different RF channels (or other frequency bands), further isolating communications associated with a particular user from communications associated with other CUGs. Such network traffic isolation improves data security at the cellular system with respect to the different CUGs (e.g., the different service providers). Additionally, flexibility of the cellular system is increased because traffic associated with a particular user is able to be processed using different rules or policies than data associated with other CUGs.

[0014] FIG. 1 illustrates an example of a system 100 configured to process network traffic associated with different virtual routing and forwarding (VRF) instances. The system 100 includes a satellite 102 that includes a channelizer 104 and a beamforming controller 108, RFE 106, a base station 118, and a core network 130 that includes a network device 142, a subscriber data server 136, and an internet protocol (IP) multimedia subsystem server 134. Although components 102, 104, 106, 108, 118, 134, 136, and 142 are illustrated as being included in the system 100, in other implementations, one or more of the components 102, 104, 106, 108, 118, 134, 136, and 142 are not included (and corresponding functions are performed one or more of the remaining components). In other implementations, the system 100 includes additional components that are not illustrated.

[0015] The satellite 102 is configured to perform communications with one or more devices (e.g., one or more subscriber devices). The one or more devices include mobile devices, such as mobile phones, tablet computers, laptop computers, smart watches, media playback devices, navigation systems, personal digital assistants (PDAs), satellite phones, vehicles (or components thereof), or a combination thereof, as non-limiting examples. In some implementations, the one or more devices also

include stationary computerized devices. The satellite 102 includes a processor 180, a memory 182, the channelizer 104, the beamforming controller 108, and a payload 184. The payload 184 includes one or more antennas configured to perform communication operations with the one or more devices. In a particular implementation, the system 100 includes or corresponds to a satellite-based cellular system. In another particular implementation, the system 100 includes or corresponds to a hybrid cellular system (e.g., a hybrid satellite and terrestrial cellular system). For example, the satellite 102 may communicate with the one or more devices via one or more components, such as antenna(s), satellite(s), tower(s), etc., of a terrestrial-based cellular system. In an alternate implementation, the system 100 includes or corresponds to a terrestrial-based cellular system, as further described with reference to FIG. 2.

[0016] The channelizer 104 is configured to divide a frequency spectrum (e.g., a RF spectrum) into multiple smaller frequency "slices" (e.g., sub-bands, channels, sub-channels, etc.). As an example, the channelizer 104 is configured to divide an incoming ultrawideband signal into channels or sub-channels for transmission to another device or component. Although RF signals are described herein, in other implementations, other frequency signals may be used. In a particular implementation, the channelizer 104 is configured to divide RF signals into narrower bandwidths than other filtering equipment. As a particular example, other filtering equipment may divide RF signals into sizes referred to as "transponders" or "transponder-size" that have a range of approximately 30-70 megahertz (MHz). In contrast, the channelizer 104 is configured to divide the RF spectrum (or another frequency spectrum) into bands or slices of approximately 31 kilohertz (kHz). Although referred to as portions of frequency bands or frequency channels (e.g., RF channels), signals may also be divided into time-based channels or subchannels (e.g., for time-division multiple access (TDMA) signals) or code-based channels or subchannels (e.g., for code-division multiple access (CDMA) signals).

[0017] The beamforming controller 108 is configured to enable beamforming to the one or more devices. In a particular implementation, the beamforming controller 108 is configured to apply beamforming weights to RF signals (e.g., "beams") to modify directivity and spatial selectivity of the RF signals to and from the one or more devices. In a particular implementation, the beamforming weights are stored as beamforming data 110 at the memory 182, and each set of beamforming weights is associated with a corresponding user (e.g., the beamforming weights are assigned on a per-user basis), as further described herein. As used herein, a "user" includes a Tier-2 operator or other service provider that provides services by using network resources of a Tier-1 operator, such as a Tier-1 operator that owns and operates the system 100. CUGs include mobile virtual network operators (MVNOs), virtual network operators (VNOs), or oth-

er service providers that provide services to subscribers using the system 100. Subscribers include people that subscribe to a service provided by the service providers (e.g., the CUGs) and access the system 100 using devices, such as a mobile phone or other equipment. In the particular implementation illustrated in FIG. 1, the channelizer 104 and the beamforming controller 108 are integrated in the satellite 102. In another particular implementation, the channelizer 104 and the beamforming controller 108 are coupled between the RFE 106 and the base station 118, as described with reference to FIG. 2.

[0018] In a particular implementation, the channelizer 104 is configured to generate a plurality of sets of RF signals associated with a corresponding portion of a RF spectrum. For example, each set of RF signals includes one or more RF signals that correspond to a particular portion of the RF spectrum, such as a particular RF channel (or a particular plurality of RF channels). The plurality of sets of RF signals include at least one RF signal associated with a particular portion of the RF spectrum that is associated with a unique user. To illustrate, the plurality of RF signals includes a first RF signal that corresponds to a particular portion of the RF spectrum (e.g., a particular RF channel) that is designated for communications associated with a first user (e.g., a first MVNO). To further illustrate, the plurality of RF signals includes a second RF signal that corresponds to a particular portion of the RF spectrum (e.g., a second particular RF channel) that is designated for communications associated with a second user (e.g., a second MVNO).

[0019] In a particular implementation, the beamforming controller 108 is configured to adjust a phase, an amplitude, or both of the RF signals based on the beamforming data, which indicates a first set of beamforming weights associated with the first user and a second set of beamforming weights associated with the second user. The RF signals are transmitted (e.g., deployed) to one or more locations to enable wireless communications with devices associated with the two service providers via different portions of the RF spectrum, as further described herein. Although illustrated as two components, in a particular implementation, the channelizer 104 and the beamforming controller 108 are integrated or correspond to a single component, or the operations of the channelizer 104 and the beamforming controller 108 are performed by processor 180.

[0020] The RFE 106 is configured to process RF signals that are received from or that are to be sent (e.g., transmitted) to the one or more devices via the satellite 102. As an example, the RFE 106 includes one or more filters, one or more amplifiers, one or more mixers, one or more other components, or a combination thereof, that process RF signals received from or that are to be sent to the one or more user equipment (UEs) via the satellite 102.

[0021] The base station 118 includes a processor 120 and a memory 122, and the base station 118 is configured to receive RF signals from the one or more devices (e.g.,

via the components 102-108) and to process the RF signals. In a particular implementation, the processor 120 is configured to execute instructions stored at the memory 122 to perform the operations described herein. In an alternate implementation, the base station 118 includes hardware, such as an application-specific integrated circuit (ASIC), a field-programmable gate array (FPGA), or a controller, that performs the operations described herein.

[0022] To illustrate, the base station 118 is configured to convert the RF signals into the digital domain to generate data packets for routing through the core network 130. To generate the data packets, the base station 118 includes multiple MODEM pools including multiple MODEMs. In a particular implementation, each MODEM pool is assigned to a particular user (e.g., a particular MVNO), and one or more MODEMs from the particular MODEM pool are assigned to process data packets associated with the corresponding user. Although described as individual MODEMs, in some implementations the base station 118 includes a single MODEM that is virtualized to provide virtual MODEM pools and virtual MODEMs that perform the operations of the MODEMs described herein. In other implementations, the MODEMs may be physical MODEMs, logical MODEMs, virtual MODEMs, or a combination thereof.

[0023] In the implementation illustrated in FIG. 1, the base station 118 includes a first MODEM 124, a second MODEM 126, and an Nth MODEM 128. Although three MODEMs are illustrated, in other implementations, N may be any integer greater than two. Each MODEM is part of a MODEM pool. As an example, a first MODEM pool 125 includes the first MODEM 124, a second MODEM pool 127 includes the second MODEM 126, and an Nth MODEM pool 129 includes the Nth MODEM 128. Although one MODEM per MODEM pool is illustrated, such illustration is not limiting. MODEM pools may include any number of MODEMs. In a particular implementation, the MODEMs of the first MODEM pool 125 are assigned to perform operations on traffic associated with a first user (e.g., a first MVNO), and the MODEMs of the second MODEM pool 127 are assigned to perform operations on traffic associated with a second user (e.g., a second MVNO).

[0024] In a particular implementation, the MODEMs are configured to generate data packets based on RF signals from devices that are associated with different CUGs. The data packets include identifiers that associate each data packet with a particular user. For example, a data packet includes an identifier of a virtual routing and forwarding (VRF) instance. The VRF instance corresponds to a user (e.g., a Tier-2 operator or other service provider) and is used to perform routing at the network device 142, as further described herein. In a particular implementation, RF signals are assigned to MODEMs for processing based on configuration data 123 stored at the memory 122. In a particular implementation, the configuration data 123 indicates an association between

device identifiers (or some other identifiers, such as session identifiers, bearer identifiers, etc.) and VRF identifiers. To illustrate, the configuration data 123 indicates that a particular device identifier (e.g., a media access control (MAC) address, an internet protocol (IP) address, etc.) is associated with a VRF instance (that corresponds to a particular user associated with a device identified by the device identifier). Data packets, including VRF instance identifiers, are sent from the base station 118 to the network device 142.

[0025] The network device 142 includes a processor 144 and a memory 146, and the network device 142 is configured to route traffic between the one or more devices and a network 156 (e.g., an external network, such as a network maintained by another Tier-1 provider, the Internet, etc.). In a particular implementation, the processor 144 is configured to execute instructions stored at the memory 146 to perform the operations described herein. In an alternate implementation, the network device 142 includes hardware, such as an ASIC, a FPGA, or a controller, that performs the operations described herein. The network device 142 may include or correspond to a signaling gateway, a packet gateway, a mobility management entity (MME), or a combination thereof. In a particular implementation, the network device 142 is a single device that performs that the functions of a signaling gateway, a packet gateway, and an MME. Alternatively, more than one communicatively coupled device may perform the operations of the network device 142.

[0026] In a particular implementation, the network device 142 is configured to receive data packets from MODEMs of the base station 118 and to generate packets to be sent through the core network 130 (e.g., the subscriber data server 136 and the IP multimedia subsystem server 134) to other devices of the network 156, as further described herein. The subscriber data server 136 (e.g., a customer charging and billing subsystem (CCBS)) is configured to monitor network use for billing and provisioning. In a particular implementation, the subscriber data server 136 is configured to provision services, track network usage, and perform billing on a per-user basis. As a non-limiting example, information regarding services and usage for subscribers is maintained and stored for each MVNO or other service provider. The IP multimedia subsystem is configured to perform operations such as initiating and maintaining voice sessions, initiating and maintaining data sessions, initiating phone rings, and other such activities that enable phone, data, and other services at the one or more devices. In some implementations, the core network 130 also includes one or more security or cybersecurity systems that are configured to provide secure logical isolation of the system 100 to external networks, such as the network 156.

[0027] During operation, the system 100 enables multiple CUGs (e.g., Tier-2 operators or service providers), such as a first user 160 (e.g., a first MVNO), a second user 164 (e.g., a second MVNO), and an Nth user 168

(e.g., an Nth MVNO) to provide network services to subscribers via the system 100. Although three service providers are illustrated, in other implementations, N may be any integer greater than two. To support multiple service providers (e.g., CUGs), a beam map and beamforming weights may be created and stored as the beamforming data 110. In a particular implementation, the beam map identifies a frequency re-use pattern for beams directed to multiple cells (e.g., earth cells) such that particular frequency bands (e.g., channels, subchannels, etc.) can be re-used in different, non-adjacent cells. In another particular implementation, the beamforming data 110 indicates a carrier map that identifies an initial capacity and carrier configuration for each beam generated by the channelizer 104. In another particular implementation, each user is associated with a corresponding cell size, beam weight, pattern, capacity, and carrier assignment, and the beamforming data 110 indicates multiple sets of beamforming weights (and other information, such as cell size, pattern, capacity, carrier assignment, etc.) associated with multiple CUGs. In some implementations, the beamforming data 110 (or portions thereof) is adjusted over time, such as based on changing subscriber density, changing traffic loading patterns, or user (e.g., MVNO) policies.

[0028] In a particular implementation, the satellite 102 transmits a plurality of RF signals (e.g., beams) in one or more patterns in one or more regions. Each pattern of RF signals corresponds to a user (e.g., a MVNO) and enables wireless communication between devices associated with the MVNO (or other service provider) and the system 100. To illustrate, the channelizer 104 receives data (e.g., from the network device 142 after performance of one or more association operations) that indicates a unique user (e.g., the first user 160) is associated with a particular portion of a frequency spectrum, and the channelizer 104 generates at least one beam associated with the particular portion. In a particular implementation, the beamforming controller 108 adjusts the plurality of RF signals based on the beamforming data 110, and each RF signal is transmitted via a portion of the frequency spectrum associated with a corresponding user. In a particular implementation, the satellite 102 stores data that indicates an association between each unique user (e.g., each MVNO) and a corresponding portion of a frequency spectrum, such as the channel mapping data 186 stored at the memory 182.

[0029] The satellite 102 may transmit (e.g., deploy) the plurality of RF signals in one or more patterns to one or more regions. In a particular implementation, the satellite 102 deploys one or more RF signals in a first pattern 162, a second pattern 166, and an Nth pattern 170. Each of the patterns 162, 166, and 170 is associated with a different user. In a particular implementation, the first pattern 162 is associated with the first user 160, the second pattern 166 is associated with the second user 164, and the Nth pattern 170 is associated with the Nth user 168. One or more of the patterns 162, 166, and 170 may be

overlaid in the same region. To illustrate, the patterns 162, 166, and 170 are overlaid in a particular region to provide wireless communication services to subscribers in the particular region via different RF channels (or other frequency bands) associated with the different CUGs (e.g., service providers such as MVNOs). Alternatively, one or more of the patterns 162, 166, and 170 may be deployed in different regions. To illustrate, a subscriber at a first region accesses the system 100 via a first RF channel, and a second subscriber at a second region accesses the system 100 via a second RF channel.

[0030] A particular example is provided to illustrate operations of the system 100. In this example, the first pattern 162, the second pattern 166, and the Nth pattern 170 are overlaid on top of each other in a first region. A first device 172 (e.g., a first subscriber device) that enters the first region synchronizes to available operator frequencies and receives data, such as master information block data and system information block data, that indicates available service providers (e.g., CUGs) that offer services via the system 100. If the first device 172 is associated with the first user 160 (e.g., if the subscriber associated with the first device 172 subscribes to a service provided by the first user), the first device 172 performs one or more association operations and performs supported wireless communications via the portion of the frequency spectrum associated with the first user 160. In the same location, a second device 174 (e.g., a second subscriber device) that is associated with the second user 164 performs one or more association operations and performs supported wireless communication via the portion of the frequency spectrum associated with the second user 164.

[0031] In this manner, the portion of the frequency spectrum associated with the first user 160 is accessible by devices (e.g., subscriber devices) located at the first location and that are associated with (e.g., subscribe to services offered by) the first user 160. The portion of the frequency spectrum is inaccessible by devices (e.g., subscriber devices) located at the first location that are not associated with the first user 160. However, devices may be associated with other CUGs and may access different portions of the frequency spectrum at the same location. Thus, different devices associated with different CUGs (e.g., different MVNOs) may access the system 100 using different frequencies at the same location.

[0032] In another particular implementation, a particular pattern of RF signals includes a first set of RF signals corresponding to a RF channel(s) associated with the first user 160, a second set of RF signals corresponding to RF channel(s) associated with the second user 164, and a third set of RF signals corresponding to RF channel(s) associated with the Nth user 168. The particular pattern is deployed at multiple different locations. As a particular example, the particular pattern is deployed at Chicago, Washington D.C., and New York. Subscribers are able to access communication services associated with each user (e.g., MVNO) using corresponding RF

channels that are the same at each location, as part of a frequency re-use design. To illustrate, subscribers located in each of the three cities are able to access services provided by the first user 160 via a first RF channel, services provided by the second user 164 via a second RF channel, and services provided by the Nth user 168 via a third RF channel. Alternatively, different patterns can be deployed to different locations. As a particular example, a first pattern including RF channels associated with the first user 160 and second RF channels associated with the second user 164 is deployed at Chicago, and a second pattern including third RF channels associated with the first user 160 and fourth RF channels associated with the Nth user 168 is deployed at New York. Subscribers associated with the first user 160 access services via the first RF channels at Chicago and via the third RF channels at New York, subscribers associated with the second user 164 access services via the second RF channels at Chicago (and do not access services at New York), and subscribers associated with the Nth user 168 access services via the fourth RF channels at New York (and do not access services at Chicago).

[0033] Performing the one or more association operations may include performing a radio link attachment operation, a network registration operation, an IP multimedia subsystem (IMS) operation, other association operations, or a combination thereof. In a particular implementation, the one or more association operations are performed by the base station 118, the network device 142, other devices, or a combination thereof. In this implementation, the base station 118 and the network device 142 store data at the memories 122, 146, respectively, that associates devices to CUGs (e.g., MVNOs) and that is updated during performance of the one or more association operations. Additional details regarding the one or more association operations are described with reference to FIG. 2.

[0034] Once devices are associated with the selected network, the devices may send RF signals to (or receive RF signals from) the system 100 via the corresponding portion of the frequency spectrum. In a particular implementation, the first device 172 sends a first RF signal via a first portion of the frequency spectrum (e.g., a first RF channel) corresponding to the first user 160, and the second device 174 sends a second RF signal via a second portion of the frequency spectrum (e.g., a second RF channel) corresponding to the second user 164. The RF signals are received by the satellite 102 and provided to the RFE 106 for processing. In a particular implementation, the RFE 106 is configured to perform one or more processing operations (e.g., amplifying, noise reduction, etc.) on the first RF signal and the second RF signal.

[0035] The base station 118 receives the first RF signal and the second RF signal and designates particular MODEMs to process the RF signals (and the resultant data packets) or incoming data packets, as further described herein. In a particular implementation, the first MODEM 124 is configured to generate a first data packet based

on the first RF signal and the second MODEM 126 is configured to generate a second data packet based on the second RF signal. The designation of particular MODEMs of particular MODEM pools to process the RF signals (or data packets) is based on the user associated with the RF signals (or data packets) and the configuration data 123. As a particular example, the configuration data 123 indicates that the first device 172 (e.g., based on an indicator of a device ID associated with the first device 172) is associated with the first user 160 that corresponds to the first VRF instance identifier and that the second device 174 is associated with the second user 164 that corresponds to the second VRF instance identifier. Based on the configuration data 123 and the indicators of the device IDs, the base station 118 (e.g., the processor 120) is configured to designate the first MODEM 124 to process the first RF signal (e.g., an RF signal or a data packet associated with the first user 160) and to designate the second MODEM 126 to process the second RF signal (e.g., an RF signal or a data packet associated with the second user 164). In a particular implementation, the indicator includes a phone number associated with a device, a MSISDN (e.g., a subscription identifier), or any other kind of indicator of a device's identity.

[0036] In a particular implementation, each MODEM of the plurality of MODEMs of the base station 118 is configured to process traffic (e.g., RF signals, data packets, etc.) associated with a corresponding VRF instance identifier. To illustrate, the first MODEM 124 and other MODEMs of the first MODEM pool 125 are configured to process RF signals and data packets associated with a first VRF instance identifier that corresponds to a first VRF instance associated with the first user 160. In a particular implementation, the first MODEM 124 is further configured to refrain from processing RF signals and data packets associated with other VRF instance identifiers, such as a second VRF instance identifier that corresponds to a VRF instance associated with the second user 164. Similarly, the second MODEM 126 and other MODEMs of the second MODEM pool 127 are configured to process data packets associated with the second VRF instance identifier and to refrain from processing RF signals and data packets associated with the first VRF instance identifier.

[0037] Processing the RF signals at the MODEMs 124-128 includes generating data packets including data and corresponding VRF instance identifiers. To illustrate, based on a determination that the first device 172 is associated with the first user 160, the base station 118 designates the first MODEM 124 to generate a first data packet based on the first RF signal. Based on a determination that the second device 174 is associated with the second user 164, the base station 118 designates the second MODEM 126 to generate a second data packet based on the second RF signal. The first data packet includes first data and the first VRF instance identifier, and the second data packet includes second data and the second VRF instance identifier.

[0038] In a particular implementation, the MODEMs 124-128 are configured to encrypt the data packets on a per-VRF (or per-user) basis. To illustrate, the first MODEM 124 is further configured to encrypt the first data packet based on a first encryption associated with the first VRF instance, and the second MODEM 126 is configured to encrypt the second data packet based on a second encryption associated with the second VRF instance. In a particular implementation, the first encryption and the second encryption are the same type of encryption, however different keys or other values are used to perform the corresponding encryptions. In an alternate implementation, the first encryption and the second encryption performed using different types of encryption. Encrypting data associated with one user (e.g., one MVNO) differently from encrypting data associated with another user (e.g., another MVNO) may improve security and separation of data corresponding to the different CUGs. To illustrate, encrypted data that is erroneously provided to a device associated with a different user is not decrypted using the same encryption as other data received by the device.

[0039] The first data packet and the second data packet are sent to the network device 142. In a particular implementation, data packets (including the first data packet and the second data packet) are sent via an Institute of Electrical and Electronics Engineers (IEEE) 802.11Q compliant connector or trunk. In a particular implementation, the connector includes a 10 gigabyte (gB) fiber pipe. In a particular implementation, the base station 118, via the MODEMs 124-128, manages the modulation and coding (MOCOD), timeslot assignment, and queues for the data packets (e.g., evolved packet system (EPS) bearers) based on one or more network quality criterions. Additional details regarding the data packets are described with reference to FIG. 2.

[0040] The network device 142 receives data packets and generates packets for transmission through the core network 130 and the network 156 to one or more other devices. For example, the network device 142 is configured to receive the first data packet and to transmit a first packet to a first device via the network 156. The network device 142 is further configured to receive the second data packet and to transmit a second packet to a second device via the network 156. In a particular implementation, the first device and the second device are the same device. In an alternate implementation, the first device and the second device are different devices. The first packet includes the first data and a first header that includes a first indicator associated with the first VRF instance, and the second packet includes the second data and a second header that includes a second indicator associated with the second VRF instance.

[0041] In a particular implementation, multiple VRF instances are maintained at the network device 142. Maintaining multiple VRF instances includes maintaining multiple routing information bases (RIBs) (e.g., routing and forwarding data) to be used for routing packets at the

network device 142. As a particular example, the network device 142 stores first VRF instance RIB data 150, second VRF instance RIB data 152, and Nth VRF instance RIB data 154 at the memory 146. Although three instances of RIB data are illustrated in FIG. 1, in other implementations, N is any other integer greater than two. The VRF instance RIB data 150-154 include or indicate routing information, forwarding information, or a combination thereof.

[0042] In a particular implementation, each of the VRF instance RIB data 150-154 is associated with a user. As an example, the first VRF instance RIB data 150 is associated with the first user 160, the second VRF instance RIB data 152 is associated with the second user 164, and the Nth VRF instance RIB data 154 is associated with the Nth user 168. The VRF instance RIB data 150-154 enables different routing paths to be established for traffic associated with different CUGs. For example, a routing path from the first device 172 to the first device of the network 156 may traverse different devices or portions of devices (e.g., different MODEMs of the base station 118) through the system 100, the core network 130, and the network 156 than a routing path from the second device 174 to the second device of the network 156. Alternatively, the routing path from the first device 172 to the first device of the network may be the same as the routing path from the second device 174 to the second device of the network.

[0043] By maintaining different RIBs for different VRF instances and by designating different resources (e.g., MODEMs) of the base station 118 for different VRF instances, logical isolation of traffic on a per-user basis is provided. To illustrate, traffic to and from the first device 172 (e.g., a subscriber device associated with the first user 160) traverses the system 100 at different timeslots or with different modulation and coding than traffic to and from the second device 174 (e.g., a subscriber device associated with the second user 164). Additionally, routes through the core network 130 and the network 156 may be different for traffic to and from the first device 172 than traffic to and from the second device 174 due to differences between the first VRF instance RIB data 150 and the second VRF instance RIB data 152. By logically isolating traffic on a per-user basis (e.g., a per-service provider or per-MVNO basis), security of a user's traffic is improved compared to systems that do not maintain different VRF instances for different CUGs.

[0044] In a particular implementation, the network device 142 stores combined RIB data 148 at the memory 146 for use in routing packets through the core network 130 and the network 156. The combined RIB data 148 is based on the VRF instance RIB data 150-154 and any additional routing information generated by the network device 142 (e.g., routing associated with packets that do not correspond to one of the VRF instances). When a data packet is received at the network device 142, the corresponding VRF instance RIB data is updated. To illustrate, the network device 142 updates the first VRF

instance RIB data 150 based on the first data packet, and the network device 142 updates the second VRF instance RIB data 152 based on the second data packet.

[0045] In a particular implementation, the network device 142 generates packets (including headers that include indicators of corresponding VRF instances) based on the received data packets and the combined RIB data 148. As an example, the network device 142 generates a first packet that includes the first data and a first header that includes a first indicator associated with the first VRF instance, and the network device 142 generates second packet includes the second data and a second header that includes a second indicator associated with the second VRF instance. The first indicator and the second indicator include or correspond to an identifier, such as a source IP address or a source port address, as non-limiting examples. In a particular implementation, the indicators correspond to labels or to "5-tuples" (e.g., a tuple that includes a source address, a source port, a destination address, a destination port, and a protocol associated with the packet). Encapsulating the first data associated with the first VRF instance using a different header than the second data associated with the second VRF instance logically isolates the first data from the second data during processing by devices of the core network 130 and the network 156.

[0046] After transmission by the network device 142, the packets traverse the core network 130 and the network 156 on a route to the eventual destination device(s). In some implementations, at least some components of the core network 130 process the packets on a per-VRF basis. In a particular implementation, the subscriber data server 136 generates (or updates) first subscriber data 138 based on the first packet and generates (or updates) second subscriber data 140 based on the second packet. The first subscriber data 138 indicates provisioning information, billing information, usage information, or a combination thereof, associated with subscribers of the first user 160, and the second subscriber data 140 indicates provisioning information, billing information, usage information, or a combination thereof, associated with subscribers of the second user 164. In a particular implementation, other components of the core network 130 process traffic on a per-VRF basis. As an example, the IP multimedia subsystem server 134 performs one or more operations on the first packet based on policies or rules corresponding to the first user 160, and the IP multimedia subsystem server 134 performs one or more operations on the second packet based on policies or rules corresponding to the second user 164. In another particular implementation, one or more security or cybersecurity devices or systems (also referred to as a "demilitarized zone" (DMZ)) processes traffic based rules or policies associated with a corresponding VRF instance. VRF instance-based rules and policies enable traffic associated with a particular user (e.g., a service provider or a Tier-2 operator) to be processed by devices of the core network 130 differently than traffic associated with other

CUGs.

[0047] Packets sent from devices of the network 156 to the subscriber devices are processed in a similar manner. The packets traverse the core network 130 and the rest of the system 100 on a per-VRF instance or per-user basis until reaching the destination subscriber device. As an example, the network device 142 is configured to receive a third packet from the first device via the network 156 and to receive a fourth packet from the second device via the network 156. The third packet includes third data and a third header that includes an indicator associated with the first VRF instance, and the fourth packet includes fourth data and a fourth indicator associated with the second VRF instance. In a particular implementation, the third packet and the fourth packet are routed to the subscriber data server 136 prior to arrival at the network device 142, and the subscriber data server 136 updates corresponding subscriber data based on the third packet and the fourth packet, as described above.

[0048] In a particular implementation, the network device 142 updates the combined RIB data 148 based on the third packet and the fourth packet if any new routing information is associated with the third packet or the fourth packet. To illustrate, the network device 142 updates the first VRF instance RIB data 150 based on the third indicator being associated with the first VRF instance, and the network device 142 updates the second VRF instance RIB data 152 based on the fourth indicator being associated with the second VRF instance. The network device 142 may also generate data packets and transmit the data packets to particular MODEMs of the base station 118 based on received packets. As an example, the network device 142 generates a third data packet that includes the third data and the first VRF instance identifier, and the network device 142 generates a fourth data packet that includes the fourth data and the second VRF identifier.

[0049] After generating the third packet and the fourth packet, the network device 142 sends the third data packet to the first MODEM 124 (e.g., based on the third indicator being associated with the first VRF instance), and the network device 142 sends the fourth data packet to the second MODEM 126 (e.g., based on the fourth indicator being associated with the second VRF instance). The base station 118 converts the third data packet and the fourth data packet into a third RF signal and a fourth RF signal, respectively, for transmission to the corresponding subscriber devices. In a particular implementation, the RF signals are processed by the components 102-108 prior to transmission to the devices. To illustrate, the RFE 106 is configured to process the third RF signal and the fourth RF signal, such as by performing filtering, amplification, etc.

[0050] The base station 118 indicates, to the RFE 106 and the satellite 102, that the different RF signals are associated with different VRF instances, and thus different CUGs (e.g., different service providers). In a particular implementation, the base station 118 associates RF

signals corresponding to different VRF instances with different logical interfaces. For example, the base station 118 associates a first RF signal that corresponds to the first VRF instance (and a second MVNO) with a first logical interface between the base station 118 and the RFE 106 (or between the base station 118 and the satellite 102), and the base station 118 associates a second RF signal that corresponds to the second VRF instance (and a second MVNO). In an alternate implementation, the base station 118 sends RF signals corresponding to different VRF instances to different ports (e.g., physical ports, logical ports, virtualized ports, etc.) of the RFE 106. As a particular example, the first MODEM 124 sends a first RF signal associated with the first VRF instance (and the first MVNO) to a first port and the second MODEM 126 sends a second RF signal associated with the second VRF instance (and the second MVNO) to a second port.

[0051] In a particular implementation, RF signals associated with different logical interfaces (or sent to different ports) are transmitted to the satellite 102 via different portions of the RF spectrum (e.g., different RF channels). As a particular example, RF signals associated with the first logical interface (and thus associated with the first VRF and the first user 160) are transmitted via a first RF channel, and RF signals associated with the second logical interface (and thus associated with the second VRF and the second user 164) are transmitted via a different RF channel. In this manner, the satellite 102 receives indication of an association between an RF signal and a VRF instance (corresponding to a Tier-2 operator or service provider)

[0052] In a particular implementation, the satellite 102 is configured to receive the RF signals from the RFE 106 and to perform beam-to-beam switching to generate RF signals to be transmitted to the one or more devices based on the channel mapping data 186. As a particular example, the processor 180 (or the channelizer 104) is configured to initiate transmission of a second RF signal from one or more antennas of the payload 184 to the first device 172 based on the channel mapping data 186 and based on a first service provider (e.g., the first user 160) associated with the first device 172. The second RF signal is based on a first RF signal received at the one or more antennas of the payload 184 from the RFE 106 (and associated with a first VRF instance). The first RF signal and the second RF signal may correspond to different RF channels. For example, the first RF signal corresponds to a first RF channel, and the second RF signal corresponds to a second RF channel that is different from the first RF channel.

[0053] In a particular implementation, the channel mapping data 186 indicates mappings of RF channels (or other portions of a frequency spectrum) for communicating with the one or more devices or the RFE 106 to service providers (e.g., CUGs). As a particular example, the channel mapping data 186 indicates that the first user 160 maps to at least the first RF channel for communi-

cations between the satellite 102 and the RFE 106 and to at least the second RF channel for communications between the satellite 102 and the first device 172. The channel mapping data 186 may be generated (or updated) by an operator of a satellite control system when service providers (e.g., MVNOs) purchase capacity and access to the system 100. For example, when the first user 160 purchases capacity and access to the system 100, the satellite control system transmits data to the satellite 102 indicating that the first user 160 is associated with the first RF channel for communications between the satellite 102 and the RFE 106 and with the second RF channel for communications between subscriber devices and the satellite 102. In some implementations, the channel mapping data 186 indicates groups of RF channels associated with each user, and at least some of each group of RF channels are associated with different locations). As additional CUGs (e.g., MVNOs) purchase capacity and access to the system 100, data indicating mappings of the additional CUGs to RF channels may be sent to the satellite 102 to update the channel mapping data 186.

[0054] In a particular implementation, the processor 180 is configured to initiate transmission of a fourth RF signal from the one or more antennas of the payload 184 to the second device 174 based on the channel mapping data 186 and based on a second service provider (e.g., the second user 164) associated with the second device 174. In this example, the fourth RF signal is based on a third RF signal received at the one or more antennas of the payload 184 from the RFE 106. The third RF signal and the fourth RF signal correspond to different RF channels, and the third and fourth RF channels correspond to different RF channels as compared to the first and second RF channels, respectively. As a particular example, the third RF signal corresponds to a third RF channel that is different from the first RF channel, and the fourth RF signal corresponds to a fourth RF channel that is different from the third RF channel and the second RF channel. In this manner, different service providers (e.g., different MVNOs) are associated with different RF channel mappings, and traffic associated with different service providers is communicated by transmission of RF signals via different RF channels. As a particular example, a RF signal that is transmitted to the first device 172 corresponds to the second RF channel (e.g., a RF channel associated with the first user 160) and a RF signal that is transmitted to the second device 174 corresponds to the fourth RF channel (e.g., a RF channel associated with the second user 164). In another particular implementation, one or more antennas of the payload 184 are configured to overlay transmission of the second RF signal and the fourth RF signal at the same location.

[0055] Additional RF signals associated with the same user may be transmitted or received via the same RF channels. As a particular example, the processor 180 is configured to initiate transmission of a sixth RF signal from the one or more antennas of the payload 184 to the first device 172 based on the channel mapping data and

based on the first service provider (e.g., the first user 160). In this example, the sixth RF signal is based on a fifth RF signal received at the one or more antennas of the payload 184 from the RFE 106. Because the fifth RF signal and the sixth RF signal are associated with the first user 160, the fifth RF signal corresponds to the first RF channel and the sixth RF signal corresponds to the second RF channel.

[0056] In a particular implementation, the satellite 102 (e.g., the processor 180, the channelizer 104, or both) send RF signals to the RFE 106 based on RF signals received from the one or more subscriber devices and the channel mapping data 186. As a particular example, one or more antennas of the payload 184 receive an RF signal via the second RF channel from the first device 172, and the processor 180 initiates transmission of an RF signal to the RFE 106 via the first RF channel. The RFE 106 receives the RF signal via the first RF channel and processes the RF signal prior to providing the RF signal to the base station 118. In a particular implementation, based on the RF signal being received via the first RF channel (e.g., a RF channel associated with the first user 160), the base station 118 assigns the RF signal to a MODEM (e.g., the first MODEM 124) of the first MODEM pool 125 for processing. Traffic received at the satellite 102 from subscriber devices associated with other CUGs (e.g., other MVNOs) is similarly processed and provided to the base station 118 via the RFE 106. In this manner, traffic sent to (or received from) subscriber devices associated with different CUGs (e.g., different MVNOs) is isolated by communications via different RF channels (or other portions of a frequency spectrum). Performing communications associated with different CUGs via different RF channels (or other portions of the frequency spectrum) may improve security and prevent data associated with a service provider (e.g., a user) from being processed together with data associated with a different service provider.

[0057] In a particular implementation, the beamforming controller 108 adjusts the RF signals based on beamforming information associated with different service providers (e.g., CUGs). As an example, the beamforming controller 108 is configured to adjust the second RF signal (e.g., the RF signal transmitted to the first device 172) based on the beamforming data 110 stored at the memory 182 and based on the first RF signal (e.g., based on the RF channel via which the first RF signal is received). In a particular implementation, the beamforming data 110 indicates sets of beamforming weights that correspond to CUGs (e.g., MVNOs) that purchase capacity and access to the system 100. As a particular example, the beamforming data 110 indicates a first set of beamforming weights associated with the first user 160 and a second set of beamforming weights associated with the second user 164.

[0058] In a particular implementation, the beamforming controller 108 adjusts RF signals corresponding to the first user 160 based on the first set of beamforming

weights and may adjust RF signals corresponding to the second user 164 based on the second set of beamforming weights. Adjusting the RF signals includes modifying an amplitude, a phase, another characteristic, or a combination thereof of the RF signals. Adjusting the RF signals based on the beamforming weights is also referred to as applying the beamforming weights to the RF signals.

[0059] In a particular implementation, the beamforming controller 108 is configured to adjust a RF signal by applying one or more first beamforming weights of the beamforming data 110 (e.g., based on the RF signal being associated with the first user 160), and the beamforming controller 108 is configured to adjust another RF signal by applying one or more second beamforming weights of the beamforming data 110 (e.g., based on the other RF signal being associated with the second user 164). Applying the beamforming weights adjusts the phase, the amplitude, another characteristic, or a combination thereof of the RF signals based on user-specific beamforming weights, which may increase directivity and spatial selectivity of the RF signals to the devices. The beamforming data 110 (or a portion thereof, such as a portion corresponding to a particular user) may be updated in accordance to changes in traffic loading patterns or user (e.g., MVNO) policy.

[0060] In a particular implementation, the channelizer 104 generates a plurality of beams based on the first RF signal and the third RF signal. To illustrate, the RF signals are received at an analog front-end of the channelizer 104, converted to the digital domain, switched through a multi-stage interconnect network, and converted back to the analog domain by the analog back-end of the channelizer 104. The RF signals are deployed by the satellite 102 via different portions of the frequency spectrum (e.g., based on the channel mapping data 186). For example, the first RF signal is converted to a particular portion of the frequency spectrum associated with the first user 160, and the second RF signal may be converted to a different particular portion of the frequency spectrum associated with the second user 164. The satellite 102 transmits the second RF signal (e.g., deploys one or more beams associated with the second RF signal in the first pattern 162) and transmits the fourth RF signal (e.g., deploys one or more beams associated with the fourth RF signal in the second pattern 166) to the first device 172 and the second device 174 via the corresponding portions of the frequency spectrum, respectively.

[0061] By maintaining different VRF instances for different CUGs (e.g., Tier-2 operators or service providers) and by communicating traffic associated with different CUGs via different RF channels (or other portions of a frequency spectrum), the system 100 may provide "end-to-end" logical isolation of traffic on a per-user basis. As an example, traffic associated with the first user 160 (e.g., the first MVNO) is encapsulated and identified based on the first VRF instance, as compared to traffic associated with the second user 164 (e.g., the second MVNO). The traffic associated with the first user 160 is routed through

the components of the system 100, such as the core network 130, and the network 156 via one or more paths that may be different than one or more paths by which traffic associated with the second user 164 (e.g., the second MVNO) is routed. To further illustrate, the traffic associated with the first user 160 is sent to (or received from) the first device 172 via different RF channels than the traffic associated with the second user 164. Providing logical isolation of traffic on a per-user basis (e.g., a per service provider basis) improves security to each service provider, as compared to processing all service providers' traffic together. Isolating traffic on a per user basis increases security of the system 100. Flexibility of the system 100 may also be increased as compared to systems that do not isolate traffic by service provider. For example, the system 100 can be configured to process traffic associated with a particular service provider (e.g., user) in accordance with different rules or policies than other service providers.

[0062] In the above description, various functions performed by the system 100 of FIG. 1 are described as being performed by the components 102-108, 118, 134, 136, and 142. This is for illustration only. In other implementations, one or more of the components 102-108, 118, 134, 136, and 142 may be integrated in a single component that performs functions associated with multiple components. Additionally, various functions are described as being performed by the processor 120 based on execution of the instructions stored at the memory 122 or by the processor 144 based on execution of the instructions stored at the memory 146. This is for illustration only. In an alternate implementation, one or more functions performed by the processor 120 or by the processor 144 are instead be performed by one or more hardware components. For example, a first component may generate packets based on the combined RIB data 148 and the first VRF instance RIB data 150. Each component may be implemented using hardware (e.g., a field-programmable gate array (FPGA) device, an application-specific integrated circuit (ASIC), a digital signal processor (DSP), a controller, etc.), software (e.g., instructions executable by a processor), or a combination thereof.

[0063] FIG. 2 illustrates an example of a system 200 configured to process network traffic associated with different VRF instances. The system 200 includes the RFE 106, the base station 118, the network device 142, the subscriber data server 136, and the IP multimedia subsystem server 134 of FIG. 1. Additionally, the system 200 includes a cellular system 202, a channelizer 204, and a beamforming controller 206. Although components 202, 204, 206, 106, 118, 134, 136, and 142 are illustrated as being included in the system 200, in other implementations, one or more of the components 202, 204, 206, 106, 118, 134, 136, and 142 are not included (and corresponding functions are performed one or more of the remaining components).

[0064] Components 106, 118, 134, 136, and 142 are configured to perform the operations described with ref-

erence to FIG. 1. The channelizer 204 and the beamforming controller 206 are configured to perform the operations described with reference to the channelizer 104 and the beamforming controller 108 of FIG. 1. As an example, the channelizer 204 generates RF signals for transmission to the one or more devices via different RF channels (or other portions of a frequency spectrum) based on channel mapping data 210, as described with reference to FIG. 1. As another example, the beamforming controller 206 adjusts one or more RF signals based on beamforming weights indicated by beamforming data 208 and associated with a particular user (e.g., a MVNO) to adjust amplitude, phase, etc., of the one or more RF signals, as described with reference to FIG. 1. However, in the implementation illustrated in FIG. 2, the channelizer 204 and the beamforming controller 206 are communicatively coupled between the RFE 106 and the base station 118 (instead of being integrated within a satellite).

[0065] In a particular implementation, the RF signals generated by the channelizer 204 (and processed by the RFE 106) are transmitted (e.g., deployed) to one or more subscriber devices by a cellular system 202. The cellular system 202 includes one or more components, such as antenna(s), satellite(s), tower(s), etc., that are configured to provide RF signals from one location to another. In a particular implementation, the cellular system 202 includes or corresponds to a terrestrial cellular system. In another particular implementation, the cellular system 202 includes or corresponds to a hybrid system (e.g., a system that includes terrestrial components, such as one or more towers, and one or more satellites).

[0066] During operation, the system 200 provides end-to-end network traffic isolation for different CUGs (e.g., Tier-2 operators or service providers), as described with reference to the system 100 of FIG. 1. However, in contrast to the system 100 of FIG. 1, channel mapping and beamforming operations are performed by terrestrial components. To illustrate, channel mapping and beamforming operations are performed by the channelizer 204 and the beamforming controller 206, respectively, which are terrestrial components and communicatively coupled between the RFE 106 and the base station 118. Reducing the number of components integrated within a satellite may decrease costs and improve ease of repair of the system 200.

[0067] FIG. 3 depicts a ladder diagram 300 that illustrates a first example of network traffic communicated via the system 100 of FIG. 1. In a particular implementation, the operations illustrated in FIG. 3 are performed by the first device 172, the second device 174, the satellite 102, the channelizer 104, the beamforming controller 108, the RFE 106, the base station 118 (including the first MODEM 124 and the second MODEM 126) and the network device 142 of FIG. 1, by the cellular system 202, the channelizer 204, and the beamforming controller 206 of FIG. 2, and by one or more devices 304. For convenience, the satellite 102, the channelizer 104, the beamforming controller 108, (or the cellular system 202, the

channelizer 204, and the beamforming controller 206) and the RFE 106 are referred to as "system front-end" 302.

[0068] As depicted in FIG. 3, the first device 172 performs first association operation(s) 310 with the base station 118 and the network device 142 to associate with the system 100 of FIG. 1 (or the system 200 of FIG. 2). The first association operation(s) 310 may include radio link attachment operations, network registration operations, IMS registration operations, or a combination thereof.

[0069] In a particular implementation, the first device 172 is associated with the first user 160 (e.g., the first device 172 subscribes to a service of the first user 160), and the first device 172 performs one or more radio link attachment operations in accordance with one or more wireless communication standards, such as a 3rd Generation Partnership Project (3GPP) standard. To illustrate, the first device 172 synchronizes to available operator frequencies and receives master information block (MIB) data and system information block (SIB) data. The MIB data and the SIB data indicate available operators associated with the system 100 (or the system 200) and information used to connect to the system 100 (or the system 200). The first device 172 initiates a random-access channel (RACH) procedure to attach to the network. As part of the RACH procedure, the first device 172 and the base station 118 perform a handshake procedure that results in establishment of a MAC-layer connection between the first device 172 and the first MODEM 124.

[0070] In a particular implementation, after the one or more radio link attachment operations, the first device 172 performs one or more network registration operations. In some implementations, the one or more network registration operations are in accordance with one or more wireless standards, such as a 3GPP standard. To illustrate, the first device 172 initiates a handshake procedure with a MME (e.g., the network device 142) to validate and authenticate the identity of the first device 172 and the network, respectively. During the handshake procedure, the MME coordinates with an equipment identity registry (EIR) to validate that the first device 172 is authorized to access the network. The MME and the EIR store and maintain data for each user (e.g., each MVNO) for performing handshake procedures. The MME manages a second handshake procedures with the SGW, the PGW, and the base station 118 that results in an allocation of a session with the core network 130 (e.g., an evolved packet core (EPC) network) and the establishment of an initial evolved packet system (EPS) bearer to and from the first device 172. The network device 142 (e.g., the MME, the SGW, the PGW) and the base station 118 are configured to store data associated with each user (e.g., each MVNO) and to support multi-VRF RIB data and queue structures. Depending on the service being requested by the first device 172, the network device 142 (e.g., the MME, the SGW, and the PGW) may create additional EPS bearers. Once the one or more

radio link attachment operations and the one or more network registration operations are complete, the first device 172 is able to transmit and receive data via the system 100 (or the system 200).

[0071] In a particular implementation, after completing the one or more network registration operations, the first device 172 performs one or more IMS registration operations. To illustrate, the first device 172 initiates an IMS registration procedure by transmitting a session initiation protocol (SIP) register messages to the IP multimedia subsystem server 134 of FIG. 1 or FIG. 2. A proxy-call setup control function (P-CSCF) of the IP multimedia subsystem server 134 receives the SIP register message and triggers handshake procedures between the P-CSCF, an interrogating CSCF (I-CSCF) of the IP multimedia subsystem server 134, and a serving CSCF (S-CSCF) of the IP multimedia subsystem server 134 that validate authorization of the first device 172 to access IMS-based services. The S-CSCF also performs a handshake procedure with a telephony application server (TAS), allowing the TAS to query a home subscriber server database for voice over long term evolution (VoLTE) data associated with the first device 172. The first device 172, the P-CSCF, and the TAS then participate in a handshake procedure to subscribe to a registration event package in order to be made aware of future changes to a registration state associated with the first device 172. The second device 174 performs second association operation(s) 312 with the base station 118 and the network device 142, in a similar manner to the first association operation(s) 310 (except that the second device 174 is associated with the second user 164). The above-described association operations are illustrative and not limiting. In other implementations, one or more other association operations are performed, one or more of the above-described association operations are not performed, or both.

[0072] After the first association operation(s) 310 are performed, the first device 172 sends a first RF signal 314 to the system front-end. The system front-end receives the first RF signal 314 and performs initial processing (e.g., amplification, filtering, noise-reduction, etc.) prior to sending the first RF signal 314 to the base station 118. The first MODEM 124 receives the first RF signal 314 and generates a first data packet 316. The first data packet includes first data 320 (e.g., data indicated by the first RF signal 314) and a first identifier 322. In a particular implementation, the first identifier 322 includes or corresponds to a label. The label includes a switching tag (also referred to as a virtual local area network (VLAN) tag or an 802.1Q tag), a VRF instance identifier, and an open shortest path first (OSPF) process identifier (PID).

[0073] The first data packet 316 is sent from the first MODEM 124 of the base station 118 to the network device 142. The network device 142 receives the first data packet 316 and generates a first packet 318 based on the first data packet 316. The first packet 318 includes the first data 320 and a first indicator 324. In a particular

implementation, the network device 142 is configured to replicate any changes made to per-VRF instance RIB data in the combined RIB data 148 (e.g., also referred to as a backbone RIB or a global VRF instance RIB). To illustrate, the combined RIB data 148 is updated based on updates to the first VRF instance RIB data 150. Additionally, the first identifier 322 is stripped from the first data 320, and the first indicator 324 is included in the first packet 318. The first indicator 324 (e.g., a "per-VRF identifier") indicates a VRF instance associated with the first user 160. In a particular implementation, the first indicator 324 is included in a header of the first packet 318. In some implementations, the first indicator 324 is indicated by a source address or a source port number of the header of the first packet 318. After generation, the first packet 318 is sent to the one or more devices 304 via the core network 130 and the network 156 of FIG. 1 or FIG. 2.

[0074] After the second association operation(s) 312 are performed, the second device 174 sends a second RF signal 330 to the system front-end. The system front-end receives the second RF signal 330 and performs initial processing (e.g., amplification, filtering, noise reduction, etc.) prior to sending the second RF signal 330 to the base station 118. The second MODEM 126 receives the second RF signal 330 and generates a second data packet 332. The second data packet includes second data 340 (e.g., data indicated by the second RF signal 330) and a second identifier 342. In a particular implementation, the second identifier 342 includes or corresponds to a label. The label includes a switching tag, a VRF instance identifier, and an OSPF PID.

[0075] The second data packet 332 is sent from the second MODEM 126 of the base station 118 to the network device 142. The network device 142 receives the second data packet 332 and generates a second packet 334 based on the second data packet 332. The second packet 334 includes the second data 340 and a second indicator 344. In a particular implementation, the network device 142 is configured to replicate any changes made to per-VRF instance RIB data in the combined RIB data 148. To illustrate, the combined RIB data 148 is updated based on updates to the second VRF instance RIB data 152. Additionally, the second identifier 342 is stripped from the second data 340, and the second indicator 344 is included in the second packet 334. The second indicator 344 indicates a VRF instance associated with the second user 164. In a particular implementation, the second indicator 344 is included in a header of the second packet 334. In some implementations, the second indicator 344 is indicated by a source address or a source port number of the header of the second packet 334. After generation, the second packet 334 is sent to the one or more devices 304 via the core network 130 and the network 156 of FIG. 1.

[0076] FIG. 4 depicts a ladder diagram 400 that illustrates a second example of network traffic communicated via the system 100 of FIG. 1 or the system 200 of FIG. 2. In a particular implementation, the operations illustrat-

ed in FIG. 4 are performed by the first device 172, the second device 174, the system front end 302 (e.g., the satellite 102, the channelizer 104, the beamforming controller 108, (or the cellular system 202, the channelizer 204, and the beamforming controller 206) and the RFE 106), the base station 118 (including the first MODEM 124 and the second MODEM 126) and the network device 142 of FIG. 1 or FIG. 2, and the one or more devices 304 of FIG. 3.

[0077] As depicted in FIG. 4, the network device 142 receives a third packet 402 from the one or more devices 304. The third packet includes third data 410 and the first indicator 324. The first indicator 324 indicates the VRF instance associated with the first user 160. In a particular implementation, the first indicator 324 is included in a header of the third packet 402. In some implementations, the first indicator 324 is indicated by a source address or a source port number of the header of the third packet 402. In a particular implementation, the network device 142 is configured to update the combined RIB data 148 based on the third packet 402 and to replicate any changes made to the combined RIB data 148 to corresponding per-VRF instance RIB data. To illustrate, the first VRF instance RIB data 150 is updated based on the updates to the combined RIB data 148.

[0078] The network device 142 receives the third packet 402 and generates a third data packet 404 based on the third packet 402. The third data packet 404 includes the third data 410 and the first identifier 322. In a particular implementation, the first identifier 322 includes or corresponds to a label. The label includes a switching tag (also referred to as a VLAN tag or an 802.1Q tag), a VRF instance identifier, and an OSPF PID. To illustrate, the first indicator 324 is stripped from the third packet 402, and the third data 410 is encapsulated in the third data packet 404 with the first identifier 322.

[0079] The first MODEM 124 receives third data packet 404 and modulates the third data 410 to generate a third RF signal 406. For example, based on the first identifier 322 identifying the first VRF instance (which corresponds to the first user 160), the third data packet 404 is received and processed at the first MODEM 124 (e.g., a MODEM of the first MODEM pool 125 which is associated with the first user 160). The third RF signal 406 represents the third data 410 and is provided by the system front end 302 for additional processing (e.g., by the RFE 106) and for transmission to the first device 172. In a particular implementation, the third RF signal 406 corresponds to a RF channel (or other portion of a frequency spectrum) that is associated with the first user 160.

[0080] The network device 142 also receives a fourth packet 420 from the one or more devices 304. The fourth packet includes fourth data 440 and the second indicator 344. The second indicator 344 indicates the VRF instance associated with the second user 164. In a particular implementation, the second indicator 344 is included in a header of the fourth packet 420. In some implementations, the second indicator 344 is indicated by a source

address or a source port number of the header of the fourth packet 420. In a particular implementation, the network device 142 is configured to update the combined RIB data 148 based on the fourth packet 420 and to replicate any changes made to the combined RIB data 148 to corresponding per-VRF instance RIB data. To illustrate, the second VRF instance RIB data 152 is updated based on the updates to the combined RIB data 148.

[0081] The network device 142 receives the fourth packet 420 and generates a fourth data packet 422 based on the fourth packet 420. The fourth data packet 422 includes the fourth data 440 and the second identifier 342. In a particular implementation, the second identifier 342 includes or corresponds to a label. The label includes a switching tag (also referred to as a VLAN tag or an 802.1Q tag), a VRF instance identifier, and an OSPF PID. To illustrate, the second indicator 344 is stripped from the fourth packet 420, and the fourth data 440 is encapsulated in the fourth data packet 422 with the second identifier 342.

[0082] The second MODEM 126 receives fourth data packet 422 and modulates the fourth data 440 to generate a fourth RF signal 424. For example, based on the second identifier 342 identifying the second VRF instance (which corresponds to the second user 164), the fourth data packet 422 is received and processed at the second MODEM 126 (e.g., a MODEM of the second MODEM pool 127 which is associated with the second user 164). The fourth RF signal 424 represents the fourth data 440 and is provided by the system front end 302 for additional processing (e.g., by the RFE 106) and for transmission to the second device 174. In a particular implementation, the fourth RF signal 424 corresponds to a RF channel (or other portion of a frequency spectrum) that is associated with the second user 164.

[0083] FIG. 5 illustrates a method 500 of processing network traffic received from a device based on a service provider (e.g., a user) associated with the device. In a particular implementation, the method 500 is performed by the system 100 of FIG. 1 or the system 200 of FIG. 2, as non-limiting examples. The method 500 is described with respect to a particular user (e.g., the first user 160) for convenience, and such description is not limiting.

[0084] The method 500 includes, after a subscriber device (e.g., user equipment (UE), radio communication terminal) associated with a MVNO has performed an association with a network, receiving a RF signal from the device via a particular RF channel that is accessible to devices associated with the MVNO, at 502. As a particular example, after the first device 172 has performed an association with the core network 130, the satellite 102 receives an RF signal from the first device 172 via an RF channel associated with the first user 160 (e.g., a first MVNO).

[0085] The method 500 includes performing front end processing on the RF signal at RFE, at 504. As a particular example, the RF signal is used to generate a second RF signal (e.g., an RF signal corresponding to a second

RF channel that is associated with the first user 160) that is transmitted to the RFE 106, and the RFE 106 performs front-end processing (e.g., amplification, noise reduction, etc.) on the second RF signal.

[0086] The method 500 includes assigning, based on the device's association with the MVNO, a first MODEM of a first MODEM pool of a base station to process the RF signal to generate a data packet that includes data and an identifier of a first VRF instance associated with the MVNO, at 506. As a particular example, the base station 118 assigns the first MODEM 124 of the first MODEM pool 125 to process the second RF signal based on the second RF signal's association with the first user 160 (as indicated by the second RF signal corresponding to the second RF channel, which is associated with the first user 160). After the second RF signal is assigned, the first MODEM 124 generates a first data packet including first data and a first VRF instance identifier that identifies a first VRF instance associated with the first user 160.

[0087] The method 500 includes generating a packet at the network device, at 508. The packet includes the data and a header including an indicator of the VRF instance based on combined RIB data. As a particular example, the network device 142 generates a first packet that includes the first data and an indicator (e.g., a label or a 5-tuple, as non-limiting examples) of the first VRF instance associated with the first user 160 based on the combined RIB data 148. The method 500 includes updating RIB data associated with the VRF instance based on the data packet, at 510, and updating the combined RIB data based on the RIB data associated with the VRF instance, at 512. As a particular example, prior to generating the packet, the network device 142 updates the first VRF instance RIB data 150 based on the first data packet, and the network device 142 updates the combined RIB data 148 based on updates to the first VRF instance RIB data 150 (or updates to any other VRF-specific RIB data).

[0088] The method 500 includes updating subscriber information associated with the first VRF instance based on the packet, at 514. As a particular example, the subscriber data server 136 updates the first subscriber data 138 (e.g., subscriber data associated with the VRF instance corresponding to the first user 160) based on the first packet.

[0089] The method 500 further includes transmitting the packet to a device via the network, at 516. As a particular example, the first packet traverses the core network 130 and the network 156 to a destination device via a first route. The first route is indicated by the first VRF instance RIB data 150, and the route may be different than a route traversed by packets associated with other VRF instances (and other CUGs).

[0090] By communicating RF signals for different CUGs (e.g., different MVNOs) via different RF channels (e.g., different portions of a frequency spectrum) and by maintaining different VRF instances for the different

CUGs, the method 500 provides end-to-end network traffic isolation for different CUGs (e.g., different Tier-2 operators or service providers). Providing end-to-end network traffic isolation reduces a likelihood that data associated with a particular user is accessible to subscribers of a different user.

[0091] FIG. 6 illustrates a method 600 of isolating processing traffic to be transmitted to a device based on a service provider associated with the device. In a particular implementation, the method 600 is performed by the system 100 of FIG. 1 or the system 200 of FIG. 2, as non-limiting examples. In some implementations, the method 600 is performed after performance of the method 500 of FIG. 5 or before performance of the method 500 of FIG. 5. The method 600 is described with respect to a particular user (e.g., the second user 164) for convenience, and such description is not limiting.

[0092] The method 600 includes receiving a packet via a network at a network device from another device, at 602. The packet includes data and a header including an indicator of a VRF instance associated with a MVNO. As a particular example, the network device 142 receives a second packet from a device via the network 156 and the core network 130. The second packet includes second data and a second header that indicates a second indicator of a second VRF instance associated with the second user 164 (e.g., a second MVNO).

[0093] The method 600 includes updating subscriber information associated with the VRF instance based on the packet, at 604. As a particular example, prior to the network device 142 receiving the second packet, the subscriber data server 136 receives the second packet and updates the second subscriber data 140 based on the second packet.

[0094] The method 600 includes generating a data packet that includes the data and an identifier of the VRF instance based on RIB data associated with the VRF instance, at 606. As a particular example, the network device 142 generates a second data packet that includes the second data and a second VRF instance identifier based on the second VRF instance RIB data 152. The method 600 includes updating combined RIB data based on the packet, at 608, and updating the RIB data associated with the VRF instance based on the combined RIB data, at 610. As a particular example, prior to generating the second data packet, the network device 142 updates the combined RIB data 148 based on the second packet, and the network device 142 updates the second VRF instance RIB data 152 based on the update to the combined RIB data 148.

[0095] The method 600 includes transmitting a data packet from the network device to a MODEM of a first MODEM pool of a base station based on the VRF instance identifier, at 612. The data packet includes the data and the VRF instance identifier. To illustrate, the base station 118 assigns the second MODEM 126 (of the second MODEM pool 127 that is associated with the second user 164) to process the second data packet

based on the second data packet including a second VRF instance identifier that corresponds to the second VRF instance (e.g., a VRF instance associated with the second user 164).

[0096] The method 600 includes generating a RF signal based on the data packet at the MODEM, at 614. As a particular example, the second MODEM 126 generates a second RF signal based on the second data packet and provides the second RF signal to the RFE 106 for front-end processing (e.g., amplification, filtering, etc.). In a particular implementation, the second MODEM 126 associates the second RF signal with a particular logical interface (or sends the RF signal to a particular port) that is associated with the second user 164.

[0097] The method 600 further includes transmitting the RF signal to a subscriber device that is associated with the MVNO via a particular RF channel, at 616. As a particular example, the satellite 102 receives the second RF signal from the RFE 106 via a particular RF channel that is associated with the second user 164. Based on the second RF signal being associated with the second user 164, the satellite 102 (or a channelizer) generates a third RF signal that corresponds to a second RF channel associated with the second user 164 and designated for communications between devices and the satellite 102 (or a channelizer). In a particular implementation, the satellite 102 transmits the third RF signal to the second device 174 via the second RF channel.

[0098] By communicating RF signals for different CUGs (e.g., different MVNOs) via different RF channels (e.g., different portions of a frequency spectrum) and by maintaining different VRF instances for the different CUGs, the method 500 provides end-to-end network traffic isolation for different CUGs (e.g., different Tier-2 operators or subscribers). Providing end-to-end network traffic isolation reduces a likelihood that data associated with a particular user is accessible to subscribers of a different user.

[0099] FIG. 7 illustrates a method 700 of processing network traffic associated with different VRF instances. In a particular implementation, the method 700 is performed by the system 100 of FIG. 1 or the system 200 of FIG. 2, such as by the network device 142, as a non-limiting example.

[0100] The method 700 includes receiving a first data packet at a network device from a base station, at 702. As a particular example, the network device includes or corresponds to the network device 142, the base station includes or corresponds to the base station 118 of FIGs. 1 and 2, and the first data packet includes or corresponds to the first data packet 316 of FIG. 3. The first data packet includes first data and a first VRF instance identifier. As a particular example, the first data includes or corresponds to the first data 320, and the first VRF instance identifier includes or corresponds to the first identifier 322.

[0101] The method 700 includes transmitting a first packet from the network device to a first device via a

network, at 704. The first packet includes the first data and a first header that includes a first indicator associated with a first VRF instance. As a particular example, the first packet includes or corresponds to the first packet 318, and the first indicator includes or corresponds to the first indicator 324 of FIG. 3.

[0102] The method 700 includes receiving a second data packet at the network device from the base station, at 706. The second data packet includes second data and a second VRF instance identifier. As a particular example, the second data packet includes or corresponds to the second data packet 332, the second data includes or corresponds to the second data 340, and the second VRF instance identifier includes or corresponds to the second identifier 342 of FIG. 3.

[0103] The method 700 further includes transmitting a second packet from the network device to a second device via the network, at 708. The second packet includes the second data and a second header that includes a second indicator associated with a second VRF instance. As a particular example, the second packet includes or corresponds to the second packet 334, and the second indicator includes or corresponds to the second indicator 344.

[0104] In a particular implementation, the method 700 further includes generating first RIB data associated with the first VRF instance based on the first data packet and generating second RIB data associated with the second VRF instance based on the second data packet. As a particular example, the first RIB data includes or corresponds to the first VRF instance RIB data 150, and the second RIB data includes or corresponds to the second VRF instance RIB data 152 of FIGs. 1 and 2. In some implementations, the method 700 further includes generating combined RIB data based on the first RIB data and the second RIB data. As a particular example, the combined RIB data includes or corresponds to the combined RIB data 148 (e.g., backbone RIB data) of FIGs. 1 and 2. In some implementations, the first header and the second header are generated based on the combined RIB data.

[0105] In another particular implementation, the method 700 further includes logically isolating data associated with the second VRF instance from data associated with the first VRF instance by encapsulating data associated with the first VRF instance using a different header than the data associated with the second VRF instance. As a particular example, data packets generated by the base station 118 and associated with different VRF instances include different VRF instance identifiers (e.g., the first identifier 322 and the second identifier 342 of FIG. 3), and packets generated by the network device 142 and associated with different VRF instances include headers that include different indicators (e.g., the first indicator 324 and the second indicator 344 of FIG. 3).

[0106] In another particular implementation, the method 700 further includes updating first subscriber information associated with the first VRF instance based on the

first packet and updating second subscriber information associated with the second VRF instance based on the second packet. As another particular example, the first subscriber information includes or corresponds to the first subscriber data 138 and the second subscriber information includes or corresponds to the second subscriber data 140 of FIGs. 1 and 2.

[0107] Although one or more of FIGS. 1-7 may illustrate systems, apparatuses, and/or methods according to the teachings of the disclosure, the disclosure is not limited to these illustrated systems, apparatuses, and/or methods. One or more functions or components of any of FIGS. 1-7 as illustrated or described herein may be combined with one or more other portions of another of FIGS. 1-7. For example, one or more elements of the method 500 of FIG. 5, one or more of the elements of the method 600 of FIG. 6, one or more of the elements of the method 700 of FIG. 7, or a combination thereof, may be performed in combination with other operations described herein. Accordingly, no single implementation described herein should be construed as limiting and implementations of the disclosure may be suitably combined without departing from the teachings of the disclosure. As an example, one or more operations described with reference to FIGS. 5-7 may be optional, may be performed at least partially concurrently, and/or may be performed in a different order than shown or described.

[0108] For example, another particular implementation comprises a system including: a base station (118) comprising: a first MODEM (124) of a first MODEM pool (125), the first MODEM (124) configured to generate a first data packet (316) based on a first radio-frequency (RF) signal (314), wherein the first data packet (316) includes first data (320) and a first virtual routing and forwarding (VRF) instance identifier (322); and a second MODEM (126) of a second MODEM pool (127), the second MODEM (126) configured to generate a second data packet (332) based on a second RF signal (330), wherein the second data packet (332) includes second data (340) and a second VRF instance identifier (342). The system further includes a network device (142) coupled to the base station (118) and configured to: receive the first data packet (316) and to transmit a first packet (318) to a first device via a network, wherein the first packet (318) includes the first data (320) and a first header that includes a first indicator (324) associated with a first VRF instance that corresponds to the first VRF instance identifier; and receive the second data packet (332) and to transmit a second packet (334) to a second device via the network, wherein the second packet (334) includes the second data (340) and a second header that includes a second indicator (344) associated with a second VRF instance that corresponds to the second VRF instance identifier.

[0109] Advantageously, the network device (142) is further configured to: receive a third packet (402) from the first device via the network, the third packet (402) including third data (410) and a third header that includes a third indicator (324) associated with the first VRF in-

stance; and transmit a third data packet (404) to the first MODEM (124) of the base station (118), the third data packet (404) including the third data (410) and the first VRF instance identifier (322). Preferably, the system further comprises a beamforming controller (108, 206) configured to adjust a third RF signal (406) based on a first set of beamforming weights (110, 208) associated with a first user, wherein the first user corresponds to the first VRF instance, and wherein the third RF signal (406) is generated based on the third data packet (404).

[0110] Advantageously, the system further comprises radio frequency equipment (RFE) (106) configured to transmit a third RF signal (406), wherein the third RF signal (406) is generated at the base station (118) based on a third data packet (404) received at the base station (118) from the network device (142), and wherein the third data packet (404) includes the first VRF instance identifier (322). Preferably, the system further comprises a channelizer (104, 204) configured to generate a plurality of sets of RF signals, each set of RF signals associated with a corresponding portion of a RF spectrum, wherein the plurality of sets of RF signals includes at least one RF signal associated with a particular portion of the RF spectrum that is associated with a unique user (160, 164, 168). Preferably, the system further comprises a cellular system (202) configured to deploy the plurality of sets of RF signals in an overlay pattern at a first location, wherein the particular portion of the RF spectrum is accessible by devices at the first location and associated with the unique user (160, 164, 168) and is inaccessible by devices that are not associated with the unique user (160, 164, 168). Preferably, the cellular system (202) comprises a satellite-based cellular system or a terrestrial-based cellular system. Preferably, the base station (118) comprises a plurality of MODEM pools (125, 127, 129) associated with a plurality of VRF instance identifiers, and wherein each MODEM (124, 126, 128) of the plurality of MODEM pools (125, 127, 129) is configured to process traffic associated with a corresponding VRF instance identifier. Preferably, the unique user (160, 164, 168) comprises a first mobile virtual network operator (MVNO) (160), and the base station (118) and the network device (142) are configured to logically isolate traffic associated with the first MVNO (160) from traffic associated with a second MVNO (164).

[0111] Advantageously, the first MODEM (124) is further configured to encrypt the first data packet (316) based on a first encryption associated with the first VRF instance, and wherein the second MODEM (126) is further configured to encrypt the second data packet (332) based on a second encryption associated with the second VRF instance.

[0112] Another particular implementation comprises a satellite (102) including: a payload (184) comprising one or more antennas; a memory (182) configured to store channel mapping data (186) that indicates a mapping of at least one radio frequency (RF) channel to a different RF channel; and a processor (180) configured to initiate

transmission of a second RF signal from the one or more antennas to a first device based on the channel mapping data (186) and based on a first service provider associated with the first device, wherein the second RF signal is based on a first RF signal received at the one or more antennas from radio frequency equipment (RFE), wherein the first RF signal corresponds to a first RF channel, and wherein the second RF signal corresponds to a second RF channel that is different from the first RF channel.

[0113] Advantageously, the processor (180) is further configured to initiate transmission of a fourth RF signal from the one or more antennas to a second device based on the channel mapping data and based on a second service provider associated with the second device, wherein the fourth RF signal is based on a third RF signal received at the one or more antennas from the RFE, wherein the third RF signal corresponds to a third RF channel that is different from the first RF channel, and wherein the fourth RF signal corresponds to a fourth RF channel that is different from the third RF channel and the second RF channel. Preferably, the one or more antennas are configured to overlay transmission of the second RF signal and the fourth RF signal at the same location.

[0114] Advantageously, the satellite (102) further comprises a beamforming controller (108) configured to adjust the second RF signal based on beamforming data (110) stored at the memory (182) and the first RF signal, wherein the beamforming data (110) indicates one or more sets of beamforming weights associated with one or more service providers.

[0115] Advantageously, the processor is further configured to initiate transmission of a sixth RF signal from the one or more antennas to the first device based on the channel mapping data (186) and based on the first service provider, wherein the sixth RF signal is based on a fifth RF signal received at the one or more antennas from the RFE, wherein the fifth RF signal corresponds to the first RF channel, and wherein the sixth RF signal corresponds to the second RF channel.

[0116] Another particular implementation comprises a method including: receiving a first data packet (316) at a network device (142) from a base station (118), the first data packet (316) including first data (320) and a first virtual routing and forwarding (VRF) instance identifier (322); transmitting a first packet (318) from the network device (142) to a first device via a network, wherein the first packet (318) includes the first data (320) and a first header that includes a first indicator (324) associated with a first VRF instance;

[0117] receiving a second data packet (332) at the network device (142) from the base station (118), the second data packet (332) including second data (340) and a second VRF instance identifier (342); and transmitting a second packet (334) from the network device (142) to a second device via the network, wherein the second packet (334) includes the second data (340) and a second header that includes a second indicator (344) associated with

a second VRF instance.

[0118] Advantageously, the method further includes: generating first routing information base (RIB) data (150) associated with the first VRF instance based on the first data packet (316); and generating second RIB data (152) associated with the second VRF instance based on the second data packet (332). Preferably, the method further comprises generating combined RIB data (148) based on the first RIB data and the second RIB data, wherein the first header and the second header are generated based on the combined RIB data (148).

[0119] Advantageously, the method further comprising logically isolating data associated with the second VRF instance from data associated with the first VRF instance by encapsulating data associated with the first VRF instance using a different header than the data associated with the second VRF instance.

[0120] Advantageously, the method further comprises: updating first subscriber information associated with the first VRF instance based on the first packet (318); and updating second subscriber information associated with the second VRF instance based on the second packet (334).

[0121] The illustrations of the examples described herein are intended to provide a general understanding of the structure of the various implementations. The illustrations are not intended to serve as a complete description of all of the elements and features of apparatus and systems that utilize the structures or methods described herein. Many other implementations may be apparent to those of skill in the art upon reviewing the disclosure. Other implementations may be utilized and derived from the disclosure, such that structural and logical substitutions and changes may be made without departing from the scope of the disclosure. For example, method operations may be performed in a different order than shown in the figures or one or more method operations may be omitted. Accordingly, the disclosure and the figures are to be regarded as illustrative rather than restrictive.

[0122] Moreover, although specific examples have been illustrated and described herein, it should be appreciated that any subsequent arrangement designed to achieve the same or similar results may be substituted for the specific implementations shown. This disclosure is intended to cover any and all subsequent adaptations or variations of various implementations. Combinations of the above implementations, and other implementations not specifically described herein, will be apparent to those of skill in the art upon reviewing the description.

[0123] The Abstract of the Disclosure is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. In addition, in the foregoing Detailed Description, various features may be grouped together or described in a single implementation for the purpose of streamlining the disclosure. Examples described above illustrate but do not limit the disclosure. It should also be understood that numerous

modifications and variations are possible in accordance with the principles of the present disclosure. As the following claims reflect, the claimed subject matter may be directed to less than all of the features of any of the disclosed examples. Accordingly, the scope of the disclosure is defined by the following claims and their equivalents.

[0124] Further examples are set out in the following clauses.

A1. A system comprising: a base station comprising: a first MODEM of a first MODEM pool, the first MODEM configured to generate a first data packet based on a first radio-frequency (RF) signal, wherein the first data packet includes first data and a first virtual routing and forwarding (VRF) instance identifier; and a second MODEM of a second MODEM pool, the second MODEM configured to generate a second data packet based on a second RF signal, wherein the second data packet includes second data and a second VRF instance identifier; and a network device coupled to the base station and configured to: receive the first data packet and to transmit a first packet to a first device via a network, wherein the first packet includes the first data and a first header that includes a first indicator associated with a first VRF instance that corresponds to the first VRF instance identifier; and receive the second data packet and to transmit a second packet to a second device via the network, wherein the second packet includes the second data and a second header that includes a second indicator associated with a second VRF instance that corresponds to the second VRF instance identifier.

A2. The system of clause A1, wherein the network device is further configured to: receive a third packet from the first device via the network, the third packet including third data and a third header that includes a third indicator associated with the first VRF instance; and transmit a third data packet to the first MODEM of the base station, the third data packet including the third data and the first VRF instance identifier.

A3. The system of clause A2, further comprising a beamforming controller configured to adjust a third RF signal based on a first set of beamforming weights associated with a first user, wherein the first user corresponds to the first VRF instance, and wherein the third RF signal is generated based on the third data packet.

A4. The system of any of clauses A1 to A3, further comprising radio frequency equipment (RFE) configured to transmit a third RF signal, wherein the third RF signal is generated at the base station based on a third data packet received at the base station from the network device, and wherein the third data packet includes the first VRF instance identifier.

A5. The system of clause A4, further comprising a

channelizer configured to generate a plurality of sets of RF signals, each set of RF signals associated with a corresponding portion of a RF spectrum, wherein the plurality of sets of RF signals includes at least one RF signal associated with a particular portion of the RF spectrum that is associated with a unique user.

A6. The system of clause A5, further comprising a cellular system configured to deploy the plurality of sets of RF signals in an overlay pattern at a first location, wherein the particular portion of the RF spectrum is accessible by devices at the first location and associated with the unique user and is inaccessible by devices that are not associated with the unique user.

A7. The system of clause A6, wherein the cellular system comprises a satellite-based cellular system or a terrestrial-based cellular system.

A8. The system of clause A6 or A7, wherein the base station comprises a plurality of MODEM pools associated with a plurality of VRF instance identifiers, and wherein each MODEM of the plurality of MODEM pools is configured to process traffic associated with a corresponding VRF instance identifier.

A9. The system of any of clauses A6 to A8, wherein the unique user comprises a first mobile virtual network operator (MVNO), and wherein the base station and the network device are configured to logically isolate traffic associated with the first MVNO from traffic associated with a second MVNO.

A10. The system of any of clauses A1 to A9, wherein the first MODEM is further configured to encrypt the first data packet based on a first encryption associated with the first VRF instance, and wherein the second MODEM is further configured to encrypt the second data packet based on a second encryption associated with the second VRF instance.

A11. A satellite comprising: a payload comprising one or more antennas; a memory configured to store channel mapping data that indicates a mapping of at least one radio frequency (RF) channel to a different RF channel; and a processor configured to initiate transmission of a second RF signal from the one or more antennas to a first device based on the channel mapping data and based on a first service provider associated with the first device, wherein the second RF signal is based on a first RF signal received at the one or more antennas from radio frequency equipment (RFE), wherein the first RF signal corresponds to a first RF channel, and wherein the second RF signal corresponds to a second RF channel that is different from the first RF channel.

A12. The satellite of clause A11, wherein the processor is further configured to initiate transmission of a fourth RF signal from the one or more antennas to a second device based on the channel mapping data and based on a second service provider associated with the second device, wherein the fourth RF signal

is based on a third RF signal received at the one or more antennas from the RFE, wherein the third RF signal corresponds to a third RF channel that is different from the first RF channel, and wherein the fourth RF signal corresponds to a fourth RF channel that is different from the third RF channel and the second RF channel.

A13. The satellite of clause A12, wherein the one or more antennas are configured to overlay transmission of the second RF signal and the fourth RF signal at the same location.

A14. The satellite of any of clauses A11 to A13, further comprising a beamforming controller configured to adjust the second RF signal based on beamforming data stored at the memory and the first RF signal, wherein the beamforming data indicates one or more sets of beamforming weights associated with one or more service providers.

A15. The satellite of any of clauses A11 to A14, wherein the processor is further configured to initiate transmission of a sixth RF signal from the one or more antennas to the first device based on the channel mapping data and based on the first service provider, wherein the sixth RF signal is based on a fifth RF signal received at the one or more antennas from the RFE, wherein the fifth RF signal corresponds to the first RF channel, and wherein the sixth RF signal corresponds to the second RF channel.

A16. A method comprising: receiving a first data packet at a network device from a base station, the first data packet including first data and a first virtual routing and forwarding (VRF) instance identifier; transmitting a first packet from the network device to a first device via a network, wherein the first packet includes the first data and a first header that includes a first indicator associated with a first VRF instance; receiving a second data packet at the network device from the base station, the second data packet including second data and a second VRF instance identifier; and transmitting a second packet from the network device to a second device via the network, wherein the second packet includes the second data and a second header that includes a second indicator associated with a second VRF instance.

A17. The method of clause A16, further comprising: generating first routing information base (RIB) data associated with the first VRF instance based on the first data packet; and generating second RIB data associated with the second VRF instance based on the second data packet.

A18. The method of clause A17, further comprising generating combined RIB data based on the first RIB data and the second RIB data, wherein the first header and the second header are generated based on the combined RIB data.

A19. The method of clause A16 or A17, further comprising logically isolating data associated with the second VRF instance from data associated with the

first VRF instance by encapsulating data associated with the first VRF instance using a different header than the data associated with the second VRF instance.

A20. The method of any of clauses A16 to A19, further comprising: updating first subscriber information associated with the first VRF instance based on the first packet; and updating second subscriber information associated with the second VRF instance based on the second packet.

Claims

1. A system comprising:

a base station (118) comprising:

a first MODEM (124) of a first MODEM pool (125), the first MODEM (124) configured to generate a first data packet (316) based on a first radio-frequency (RF) signal (314), wherein the first data packet (316) includes first data (320) and a first virtual routing and forwarding (VRF) instance identifier (322); and

a second MODEM (126), the second MODEM (126) being part of a second MODEM pool (127) and configured to generate a second data packet (332) based on a second RF signal (330), wherein the second data packet (332) includes second data (340) and a second VRF instance identifier (342); and

a network device (142) coupled to the base station (118) and configured to:

receive the first data packet (316) and to transmit a first packet (318) to a first device via a network, wherein the first packet (318) includes the first data (320) and a first header that includes a first indicator (324) associated with a first VRF instance that corresponds to the first VRF instance identifier; and

receive the second data packet (332) and to transmit a second packet (334) to a second device via the network, wherein the second packet (334) includes the second data (340) and a second header that includes a second indicator (344) associated with a second VRF instance that corresponds to the second VRF instance identifier.

2. The system of claim 1, wherein the network device (142) is further configured to:

receive a third packet (402) from the first device via the network, the third packet (402) including third data (410) and a third header that includes a third indicator (324) associated with the first VRF instance; and

transmit a third data packet (404) to the first MODEM (124) of the base station (118), the third data packet (404) including the third data (410) and the first VRF instance identifier (322).

3. The system of claim 2, further comprising a beam-forming controller (108, 206) configured to adjust a third RF signal (406) based on a first set of beam-forming weights (110, 208) associated with a first user, wherein the first user corresponds to the first VRF instance, and wherein the third RF signal (406) is generated based on the third data packet (404).

4. The system of any of claims 1 to 3, further comprising radio frequency equipment (RFE) (106) configured to transmit a third RF signal (406), wherein the third RF signal (406) is generated at the base station (118) based on a third data packet (404) received at the base station (118) from the network device (142), and wherein the third data packet (404) includes the first VRF instance identifier (322).

5. The system of claim 4, further comprising a channelizer (104, 204) configured to generate a plurality of sets of RF signals, each set of RF signals associated with a corresponding portion of a RF spectrum, wherein the plurality of sets of RF signals includes at least one RF signal associated with a particular portion of the RF spectrum that is associated with a unique user (160, 164, 168).

6. The system of claim 5, further comprising a cellular system (202) configured to deploy the plurality of sets of RF signals in an overlay pattern at a first location, wherein the particular portion of the RF spectrum is accessible by devices at the first location and associated with the unique user (160, 164, 168) and is inaccessible by devices that are not associated with the unique user (160, 164, 168).

7. The system of claim 6, wherein the base station (118) comprises a plurality of MODEM pools (125, 127, 129) associated with a plurality of VRF instance identifiers, and wherein each MODEM (124, 126, 128) of the plurality of MODEM pools (125, 127, 129) is configured to process traffic associated with a corresponding VRF instance identifier.

8. The system of claim 6 or claim 7, wherein the unique user (160, 164, 168) comprises a first mobile virtual network operator (MVNO) (160), and wherein the base station (118) and the network device (142) are configured to logically isolate traffic associated with

the first MVNO (160) from traffic associated with a second MVNO (164).

9. A satellite (102) comprising:

a payload (184) comprising one or more antennas;
a memory (182) configured to store channel mapping data (186) that indicates a mapping of at least one radio frequency (RF) channel to a different RF channel; and
a processor (180) configured to initiate transmission of a second RF signal from the one or more antennas to a first device based on the channel mapping data (186) and based on a first service provider associated with the first device, wherein the second RF signal is based on a first RF signal received at the one or more antennas from radio frequency equipment (RFE), wherein the first RF signal corresponds to a first RF channel, and wherein the second RF signal corresponds to a second RF channel that is different from the first RF channel.

10. The satellite (102) of claim 9, wherein the processor (180) is further configured to initiate transmission of a fourth RF signal from the one or more antennas to a second device based on the channel mapping data and based on a second service provider associated with the second device, wherein the fourth RF signal is based on a third RF signal received at the one or more antennas from the RFE, wherein the third RF signal corresponds to a third RF channel that is different from the first RF channel, and wherein the fourth RF signal corresponds to a fourth RF channel that is different from the third RF channel and the second RF channel.

11. The satellite (102) of claim 10, wherein the one or more antennas are configured to overlay transmission of the second RF signal and the fourth RF signal at the same location.

12. The satellite (102) of any of claims 9 to 11, further comprising a beamforming controller (108) configured to adjust the second RF signal based on beamforming data (110) stored at the memory (182) and the first RF signal, wherein the beamforming data (110) indicates one or more sets of beamforming weights associated with one or more service providers.

13. A method comprising:

receiving a first data packet (316) at a network device (142) from a base station (118), the first data packet (316) including first data (320) and a first virtual routing and forwarding (VRF) in-

stance identifier (322);

transmitting a first packet (318) from the network device (142) to a first device via a network, wherein the first packet (318) includes the first data (320) and a first header that includes a first indicator (324) associated with a first VRF instance;

receiving a second data packet (332) at the network device (142) from the base station (118), the second data packet (332) including second data (340) and a second VRF instance identifier (342); and

transmitting a second packet (334) from the network device (142) to a second device via the network, wherein the second packet (334) includes the second data (340) and a second header that includes a second indicator (344) associated with a second VRF instance.

14. The method of claim 13, further comprising:

generating first routing information base (RIB) data (150) associated with the first VRF instance based on the first data packet (316); and
generating second RIB data (152) associated with the second VRF instance based on the second data packet (332).

15. The method of claim 14, further comprising generating combined RIB data (148) based on the first RIB data and the second RIB data, wherein the first header and the second header are generated based on the combined RIB data (148).




FIG. 1




FIG. 2




FIG. 3


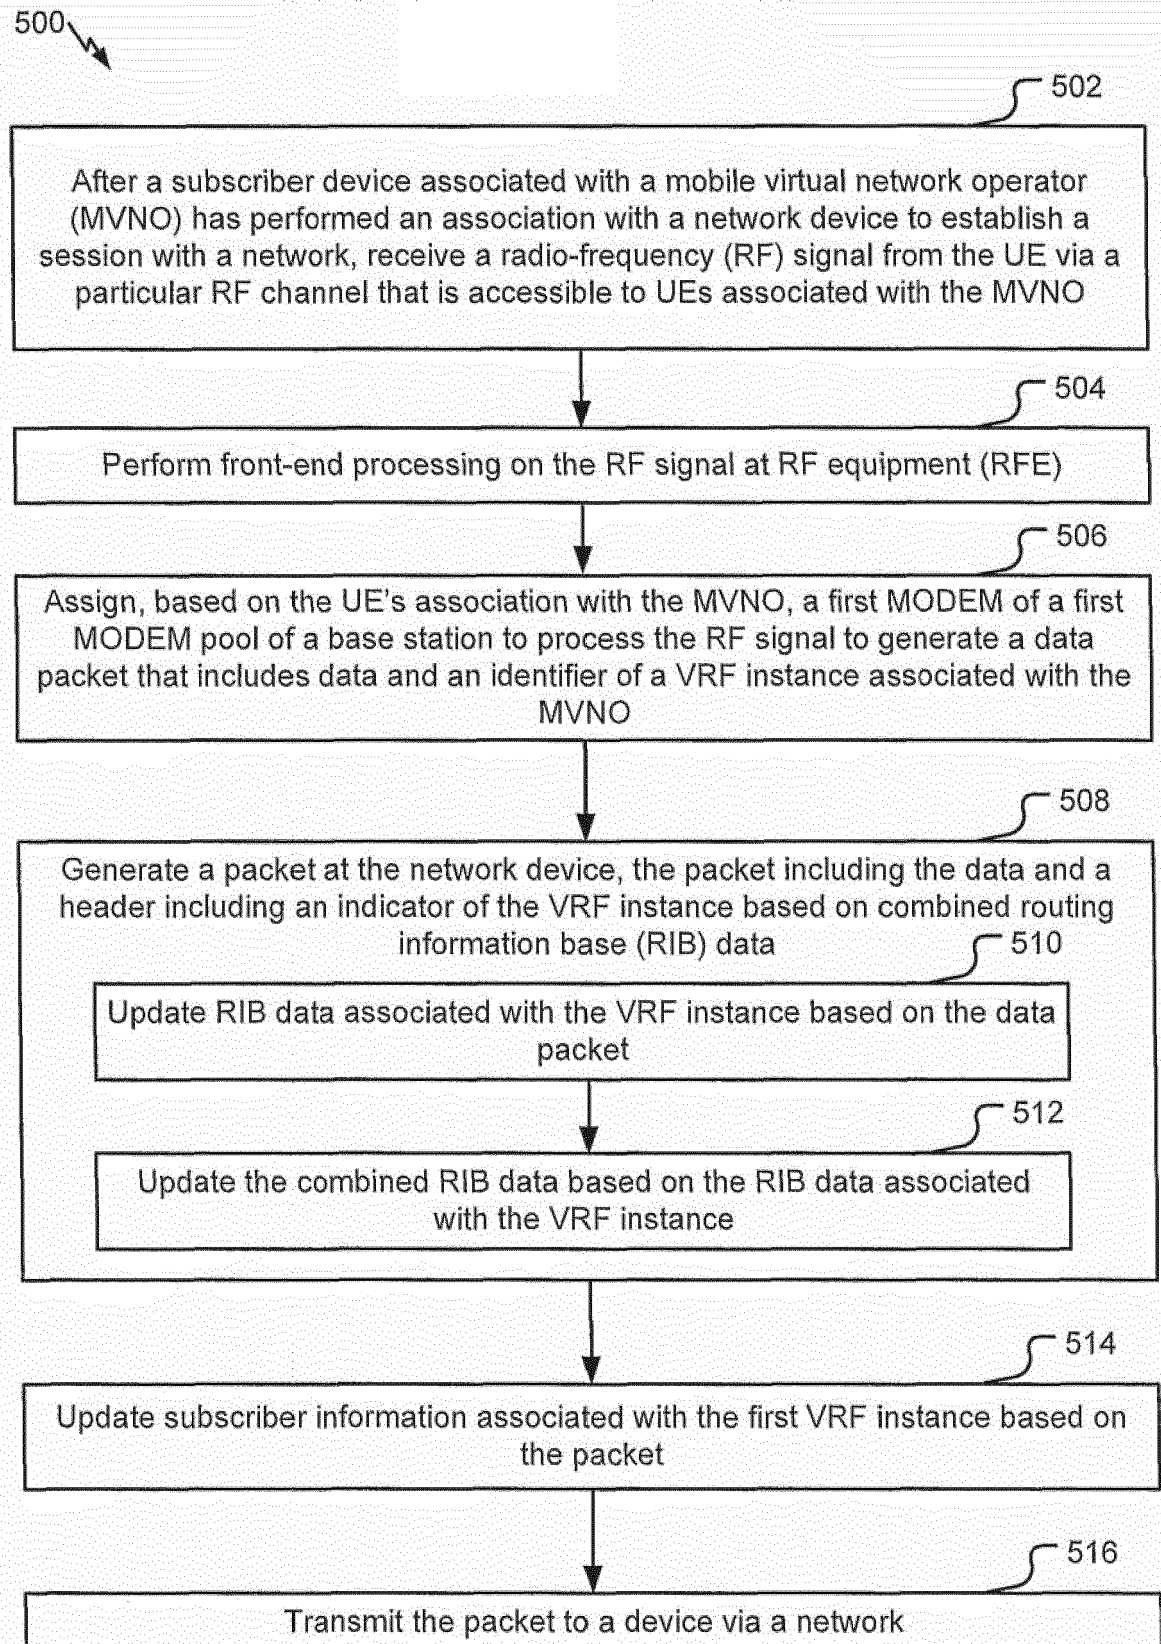
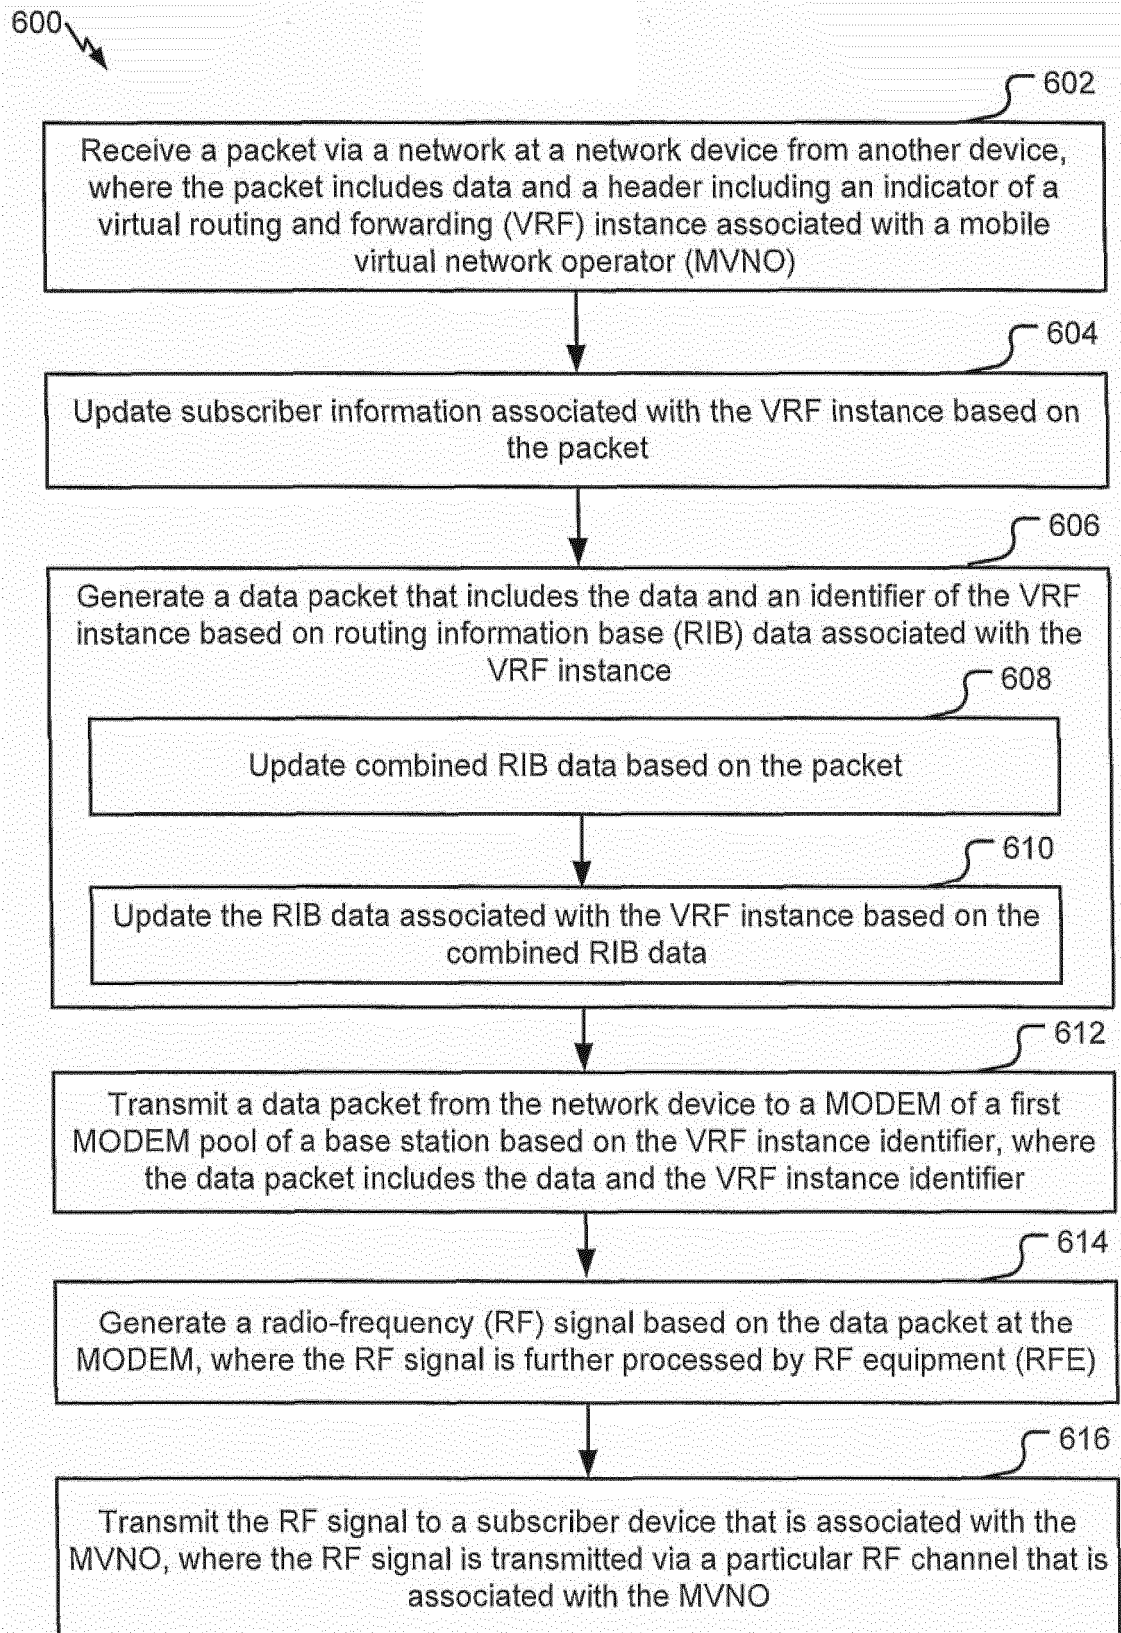
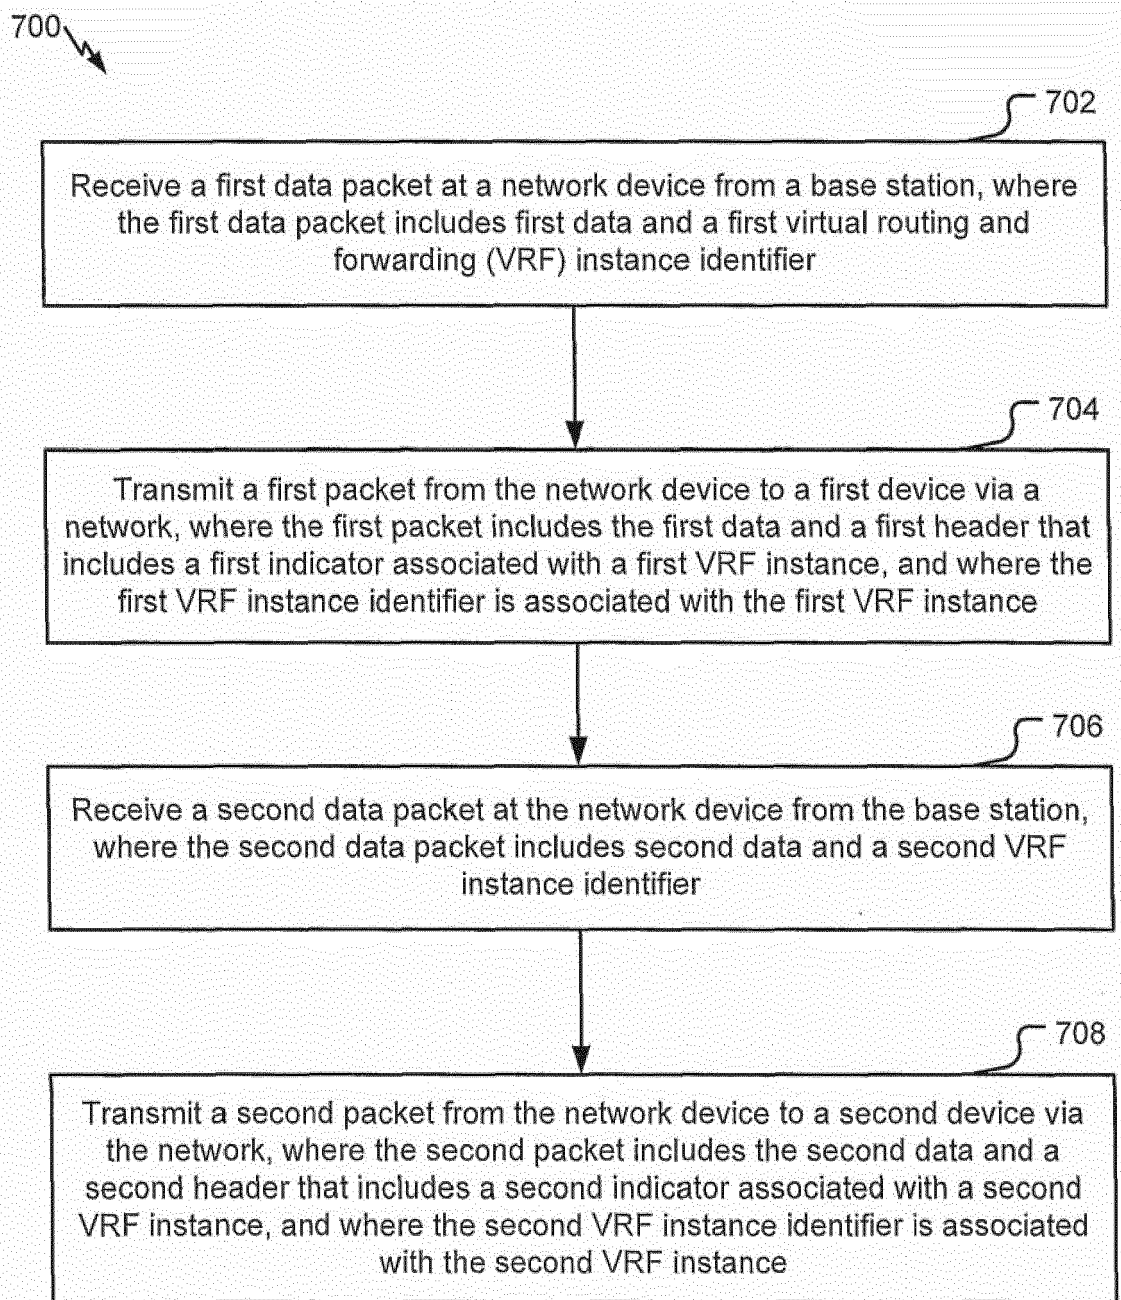


FIG. 4

**FIG. 5**

**FIG. 6**

**FIG. 7**